NBER WORKING PAPER SERIES

RATIONAL INATTENTION AND SEQUENTIAL INFORMATION SAMPLING

Benjamin Hébert
Michael Woodford

Working Paper 23787
http://www.nber.org/papers/w23787

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2017

The authors would like to thank Mira Frick, Matthew Gentzkow, Emir Kamenica, Divya Kirti, Ran Shorrer, Stephen Morris, Ming Yang, Doron Ravid, Mark Dean, Pietro Ortoleva, Sebastian Di Tella, Ilya Segal, Jacob Leshno, and participants at the Cowles Theory conference, Barcelona GSE Summer Conference on Stochastic Choice, and Stanford GSB research lunch for helpful discussions on this topic, and the NSF for research support. Portions of this paper 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.

Rational Inattention and Sequential Information Sampling
Benjamin Hébert and Michael Woodford
NBER Working Paper No. 23787
September 2017
JEL No. D83,E70

ABSTRACT

We propose a new principle for measuring the cost of information structures in rational inattention problems, based on the cost of generating the information used to make a decision through a dynamic evidence accumulation process. We introduce a continuous-time model of sequential information sampling, and show that, in a broad class of cases, the choice frequencies resulting from optimal information accumulation are the same as those implied by a static rational inattention problem with a particular static information-cost function. Among the static cost functions that can be justified in this way is the mutual information cost function proposed by Sims (2010), but we show that other cost functions can be micro-founded in this way as well. In particular, we introduce a class of "neighborhood-based" cost functions, which make it more costly to undertake experiments that can produce different results in similar states, and show that the predictions of this alternative rational inattention theory better conform with evidence from perceptual discrimination experiments.

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

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

An online appendix is available at http://www.nber.org/data-appendix/w23787

1 Introduction

The theory of rational inattention, proposed by Christopher Sims and surveyed in Sims (2010), endogenizes the imperfect awareness that decision makers have about the circumstances under which they must choose their actions. According to the theory, a decision maker (DM) chooses her action on the basis of a subjective representation of the decision situation that can be treated (formally) as a signal transmitted to the DM by a sender with knowledge of the true state. The signal structure is assumed to be optimal, in the sense of allowing the best possible state-contingent action choice, subject to a cost of more informative signal structures. In Sims' theory, the cost of an arbitrary signal structure is proportional to the Shannon's mutual information between the state of the world (that determines the DM's reward from choosing different actions) and the signal received by the DM.

In this paper, we will introduce the "neighborhood-based cost functions" as an alternative to Shannon's mutual information for use in rational inattention problems. We demonstrate that these cost functions have three appealing features. First, they satisfies the standard conditions (e.g. the conditions of De Oliveira et al. (2017) and Caplin and Dean (2015)) typically assumed for cost functions in rational inattention problems. Second, when used in a static rational inattention problem, they summarize the choice probabilities that arise from a dynamic problem, in which the DM gathers information sequentially before making her choice. Third, the neighborhood-based cost functions can capture the notion that certain pairs of states are "similar," meaning that it is difficult for the DM to discriminate between these states. Shannon's mutual information has the first two of these features, but not the third; we will provide an example, motivated by perceptual experiments, in which rationally inattentive DMs' behavior is consistent with the neighborhood-based cost function, but not the mutual-information cost function. We also argue that this distinction has consequences in economic applications, focusing on the global games example of Morris and Yang (2016). We provide an example neighborhood-based cost function, the expected Fisher information, which, like mutual information, has only a single parameter, but unlike mutual information, captures the notion that similar states are difficult to distinguish.

Sims' original work on rational inattention motivated the use of Shannon's mutual information, a measure of the degree to which the signal is informative about the state of the world, by referencing the central role it plays in information theory (Cover and Thomas (2012)). Mutual information is central to information theory as a consequence of powerful mathematical results that are of considerable practical relevance in communications engineering. It is not obvious, though, that the theorems that justify the use of mutual information in communications engineering 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 in Woodford (2012). For example, the mutual-information cost function imposes a type of symmetry across different states of nature, so that it is equally easy or difficult to distinguish between any two states that are equally probably ex ante. In the experimental task discussed by Caplin and Dean (2015), in which subjects are presented with an array of 100 red and blue balls, and must determine whether there are more red balls or more blue on a given trial, Sims' theory of rational inattention implies that, be-

[1] As 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 the kind of situation with which we are concerned, an action must be taken in an individual instance of a decision problem, and we wish to formulate a constraint on the precision of the subjective perception of the decision problem in the individual instance, without supposing that subjective descriptions of a long sequence of similar problems can be jointly formed before deciding how to act in any of the individual problems.

cause the reward from any action (e.g., declaring that there are more red balls) is the same for all states with the property that there are more red balls than blue, the probability of a subject's choosing that response will be the same in each of those states. In fact, it is much easier to quickly and reliably determine that there are more red balls for some arrays in this class (e.g. one with 98 red balls and only two blue balls) than others (e.g. one with 51 red balls and 49 blue balls, relatively uniformly dispersed), and subjects make more correct responses in the former case.[2]

One response to unappealing features of the mutual-information cost function is to develop a theory of rational inattention that makes only much weaker assumptions about the cost function (typically assumptions that are consistent with mutual information, but not requiring it), as authors such as De Oliveira et al. (2017), Caplin and Dean (2015), and Huettner et al. (2016) have done. This approach results in a theory with correspondingly weaker predictions. We seek instead to motivate a more specific class of information-cost functions, so as to allow more definite conclusions, while still including cases that we regard as more realistic specifications than mutual information.

Our approach exploits the special structure implied by an assumption that information sampling occurs through a sequential process, in which each additional signal that is received determines whether additional information will be sampled, and if so, the kind of experiment to be performed next. We emphasize the limiting case in which each individual experiment is only minimally informative, but a very large number of independent experiments can be performed. In this continuous-time limit, we obtain strong and relatively simple characterizations of the implications of rational inattention, owing to the fact that only local properties of the assumed cost function for individual experiments matter in this

[2]Dewan and Neligh (2017) present results from a related experiment in which subjects must instead estimate the number of dots in a visual array. In this case, Sims' theory would predict that all erroneous estimates should occur with equal frequency, since the reward for an erroneous response is the same regardless of which erroneous response is given. Instead, subjects are more likely to offer incorrect estimates that are close to the correct number than incorrect estimates that are far from the correct number.

limiting case.

We believe that it is often quite realistic to assume that information is acquired through a sequential sampling process. As discussed in Fehr and Rangel (2011) and Woodford (2014), an extensive 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 perceptual classification based on sequential sampling. More recently, some authors have proposed that data on stochastic choice and response time in economic contexts can be similarly modeled.[3]

Much of the empirical literature that models stochastic choice as the outcome of a sequential sampling process aims simply to provide a process model of observed behavior, but a number of recent papers endogenize at least some aspects of the information-sampling process. Fudenberg et al. (2015) consider the optimal stopping problem for an information sampling process described by the sample path of a Brownian motion with a drift that depends on the unknown state of the world.[4] This can be thought of as a problem in which a given experiment (with the probability of positive or negative outcomes dependent on the state of the world) can be repeated an indefinite number of times, with a fixed cost per repetition of the experiment. The sequence of outcomes of the successive experiments becomes a Brownian motion in the limiting case in which individual experiments require only an infinitesimal amount of time (and hence involve only an infinitesimal cost, as a fixed cost of sampling per unit time is assumed), and are correspondingly minimal in the information that they reveal about the state (because the difference in the mean outcome of an individual experiment across different states of the world is tiny relative to the standard

[3] In addition to the references in Fehr and Rangel (2011), recent examples include Krajbich et al. (2014) and Clithero (2016). In the case of preferential choice between goods, the sampling presumably does not refer to a sequence of repeated observations that are required to verify that the candy bar offered for sale is indeed a Snickers bar, but more plausibly to a sequence of repeated value assessments used to estimate the value of this option to the DM. These might be draws from memory (associations or recollections of past experiences), as suggested by Shadlen and Shohamy (2016).

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

deviation of the outcome). In the case assumed by Fudenberg et al. (2015), there is no choice about the type of experiment that can be repeatedly performed, but the decision when to stop collecting further information is optimized.

Woodford (2014) instead takes as given a stopping rule (motivated by the empirical psychology and neuroscience literatures), but endogenizes the information sampling process, as in theories of rational inattention. The assumed stopping rule makes the decision whether to continue sampling (and the event action chosen) a function only of a single number, the cumulative excess of positive over negative signals from the sequence of experiments; but at each point in the sequential process, it is assumed to be possible to vary the probability of a positive response conditional upon the true state, subject to an information-cost function that makes more informative experiments more costly. Under the assumed cost function, Woodford (2014) finds that optimal information sampling results in an evidence process that evolves as a continuous stochastic diffusion process, and that it is optimal for the drift of this process to be an increasing function of the relative value of the two choice options, as assumed by Fudenberg et al. (2015); but rather than this process being a Brownian motion with a constant drift, as in Fudenberg et al. (2015), it is generally optimal for the drift also to depend on the current belief state.

Our approach differs from these earlier efforts in seeking to endogenize *both* the nature of the information that is accumulated at each stage of the information-sampling process and the stopping rule that determines how much information is collected before a decision is made. We also consider decision problems with an arbitrary finite number of choice alternatives, rather than restricting attention to binary choice problems, as in both Fudenberg et al. (2015) and Woodford (2014). In the sequential information sampling problem considered here, we allow the information sampled at each stage to be chosen very flexibly, as in Woodford (2014), subject only to a "flow" information-cost function; but we also allow the decision when to stop sampling and make a decision to be made optimally, on the basis of the entire history of information sampled to that point, as in Fudenberg et al. (2015). Among other results, we describe a class of information-cost functions such that in the case of a binary decision, the DM's beliefs evolve according to a diffusion along a one-dimensional line segment, with a decision being made when either of the two endpoints is reached, as postulated by Woodford (2014).

In our continuous time model, the optimal information-sampling problem is presented as a problem of optimal control of a diffusion process on the probability simplex (the set of possible posterior beliefs), with sampling stopping when certain (endogenous determined) boundaries are reached. For a special (but still relatively flexible) family of possible cost functions for individual experiments, the continuous time model's predictions with regard to choice frequencies conditional on the state of the world are the same as those of a static rational-inattention model, with an appropriately chosen information-cost function for the choice of a single signal. The finite set of possible signals in the equivalent static model corresponds to the set of different possible *terminal* information states in the dynamic model, each of which corresponds to one of the possible actions. For a particular family of flow information-cost functions, the cost function for the equivalent static model is just the mutual information between the action chosen and the true state of the world; we thus provide foundations for the kind of rational inattention problem proposed by Sims (2010), that do not rely on any analogy with rate-distortion theory in communications engineering.

While our dynamic model makes predictions that are equivalent to those of the rational inattention theory of Sims (2010) (and more particularly, its application to stochastic choice by Matêjka et al. (2015)) for this particular family of flow information-cost functions, we show that different predictions can be obtained under other, very plausible specifications of the flow cost function. We focus on the implications of an attractive family of flow information-cost functions, which we call "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 cost functions of this kind from the mutual-information cost function assumed by Sims. We show that versions of our theory that assume a flow information-cost function in this family can explain the kind of continuous variation of response frequencies with changes in the characteristics of the alternatives presented that is commonly observed in perceptual discrimination experiments (but that would not be predicted by the standard theory of rational inattention).

As a still more specific special case, we consider a neighborhood-based cost function that can be defined in the limiting case of a continuum of states that can be ordered on a line — a case of considerable interest, both for economic applications and for applications to perceptual psychology. We consider the continuous-state limit of a static rational inattention problem with a particular type of neighborhood-based cost function, and show that it is equivalent to a static rational inattention problem with a particular cost function, proportional to the average Fisher information (rather than the mutual information) of the continuous family of statistical models defined by the information structure. Like Sims’ proposal, this proposal yields a highly parsimonious theory of rational inattention with only one free parameter, indexing the degree of overall scarcity of attentional resources, but one that we believe may be preferable to Sims’ proposal for many applications, such as the problem of endogenous information in global games discussed by Morris and Yang (2016).

Our paper builds upon the rational inattention literature, surveyed in Sims (2010). In its use of axioms to characterize the assumed form of the flow information-cost function, it is particularly close to Caplin and Dean (2015), Caplin et al. (2017), and De Oliveira et al. (2017). The Chentsov (1982) theorems used to characterize the properties of general rational inattention cost functions were also used by Hébert (2014), in a different context.

We also use techniques developed by Kamenica and Gentzkow (2011) and Matêjka et al. (2015) in characterizing the solution to our problem.

Section 2 begins directly with a description of our continuous-time model, and introduces the information-cost matrix function as a way of parameterizing information costs in this model. Section 3 then presents one of our main results (Theorem 1), that in a large set of cases, the solution to the continuous-time model is equivalent, in terms of the joint distribution of choices and states, to the implications of a static rational inattention model with a suitable static information-cost function. In section 4, we discuss the connection between the information-cost matrix function of the continuous-time model and the flow information-cost function for an individual signal, and state a set of general assumptions that flow information-cost functions are assumed to satisfy, in the spirit of the treatment of static rational inattention problems by De Oliveira et al. (2017).

In section 5, we introduce a specific class of flow cost functions that satisfy these general conditions, and are of interest because they incorporate a notion of "distance" between different states of the world; we then derive the static information-cost functions that correspond to them, as possible alternatives to mutual information as a cost function in rational inattention models. In subsection 5.2 of this section, we propose an even more specific static cost function, based on the Fisher information as a measure of the informativeness of an information structure, that can be applied to rational inattention models with a continuum of states.

Finally, section 6 provides a justification for the continuous-time model proposed in section 2, and for the connection between flow cost function specifications and information-cost matrix function of the continuous-time model asserted in section 4. Here we show that a discrete-time dynamic evidence accumulation problem, in which the cost of each individual signal is given by a flow cost function satisfying the assumptions stated in section 4, leads to the continuous-time problem discussed in section 2, in the limit as the number of successive signals per time period is made large, while the informativeness of each individual signal is made small at a corresponding rate. Section 7 concludes.

2 A Continuous-Time Model of Sequential Evidence Accumulation

We begin by directly introducing our continuous-time model of sequential evidence accumulation, leaving for later (section 6) the demonstration that it arises as a limiting case of an explicit discrete-time dynamic evidence accumulation problem. Let $x \in X$ be the underlying state of the nature, and $a \in A$ be the action taken by the DM. For simplicity, we assume that A and X are finite sets. We also assume that 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. We assume a penalty of this kind, rather than time discounting, for reasons of tractability.

The DM does not perfectly observe the state $x \in X$. At each time t, the DM holds beliefs $q_t \in \mathscr{P}(X)$, where $\mathscr{P}(X) \subset \mathbb{R}^{|X|}$ 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 simply choose a to maximize $u_a^T \cdot q_\tau$, where u_a is the vector of utilities associated with action a, resulting in payoff $\hat{u}(q_\tau) - \kappa\tau$.

When the DM gathers information, she chooses the variance-covariance matrix of possible changes in her beliefs, subject to certain constraints. In our model, 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 $|X|-1$-dimensional Brownian motion,[5] σ_t is a matrix that can be chosen by the DM, and $\sigma_{x,t}$ is a particular row of that matrix.

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

$$\iota^T \cdot dq_t = 0$$

always, where ι 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. Our notation enforces the requirement that $dq_{x,t} = 0$ if $q_{x,t} = 0$.

The non-trivial restriction, which limits the quantity of information the DM can acquire at each moment, is

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

where $k(q_t)$ is an $|X| \times |X|$ dimensional matrix-valued function we will refer to as the "information-cost matrix function", $tr[\cdot]$ 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. For now, we note simply that the information-cost matrix function satisfies certain properties: for any q_t, $k(q_t)$ is symmetric and positive

[5] Note that this is largest possible number of independent Brownian motions of which dq_t may be a linear combination.

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.[6]

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 constraints listed previously.

Wherever this value function is twice-differentiable and the DM does not choose to stop, the problem can be given a simple recursive representation:

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

subject to the information constraint (2), where $D(q_t)$ 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 $q = q_t$.[7]

The following lemma describes the Hamilton-Jacobi-Bellman (HJB) equation associated with this dynamic optimization problem. It is derived by showing that the information constraint binds.[8] The maximum eigenvalue appears in place of a maximization over σ_t,

[6] Actually, because we require that $\sigma_t \in M(q_t)$, constraint (2) only involves the quadratic form $v^T k(q_t) w$ defined for vectors v and w such that $v^T q_t = w^T q_t = 0$. We extend the definition of the quadratic form to all vectors $v, w \in \mathbb{R}^{|X|}$, in order to obtain a unique representation in terms of a matrix $k(q_t)$, by adding the requirement that $k(q)v = 0$ for any vector $v \in \mathbb{R}^{|X|}$ with the property that v_x is equal to a constant for all x in the support of q.

[7] In the case of a differentiable function $V(q)$ defined on the probability simplex $\mathscr{P}(X)$, in order to write the Hessian of the function as a matrix, we must adopt a coordinate system for the tangent space to the probability simplex. Throughout this paper, we do this by extending the function to the domain $\mathbb{R}_+^{|X|}$ by defining the function to be homogeneous of degree one on this larger domain (an assumption that does not restrict the function's values on the simplex). Vectors in the tangent space are then 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), (11), (15), and (17) below should also all be understood in this way.

[8] The derivation depends on an additional property of the $k(q_t)$ matrix that will be discussed below.

but this is just a compact way of expressing the idea that the DM is choosing in which direction(s) to update her beliefs.

Lemma 1. *Anywhere the value function $V(q_t)$ is twice-differentiable, it satisfies*

$$\max\{\lambda_1(D(q_t)V_{qq}(q_t)D(q_t) - \theta k(q_t)), \hat{u}(q_t) - V(q_t)\} = 0, \tag{3}$$

where $\theta = \chi^{-1}\kappa$, and for any $|X| \times |X|$ matrix K, $\lambda_1(K)$ denotes the largest eigenvalue of K associated with an eigenvector v such that $\iota^T v = 0$.[9]

Proof. See the appendix, section A.1. □

This equation has the standard form of an optimal stopping problem, with the twist that it is a "Hessian equation" in the continuation region. The parameter θ 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. The caveat about twice-differentiability plays several roles. First, as is common in optimal stopping problems, the value function may not be twice differentiable on the stopping boundary. Second, the Hessian equation in the continuation region is "degenerate elliptic", and therefore a solution that is twice-differentiable everywhere in the continuation region may not exist. A third complication is that the beliefs q_t may come to place zero weight on a certain state — that is, the beliefs may hit the boundary of the simplex, at which point the value function $V(q_t)$ is not twice-differentiable in all directions. Fortunately, in what follows, these issues will be a nuisance, rather than a serious obstacle.

[9]Here we are interested in the eigenvectors of the matrix corresponding to elements of the tangent space to the probability simplex. Note that under our notation for writing quadratic forms over the probability simplex as matrices, explained in footnotes 6 and 7 above, ι is a null eigenvector of both $D(q)V_{qq}D(q)$ and $k(q)$, for any q; but we do not wish to count this as one of the eigenvectors of the linear operator for purposes of defining the maximum eigenvalue, as our first-order condition actually involves a linear operator defined on the tangent space of the probability simplex.

The DM's optimal stopping rule is characterized by the standard value-matching and smooth-pasting conditions. Let $\Omega \subset \mathscr{P}(X)$ 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.[10] However, it will never be optimal for the DM to stop at one of these indifference points.

Before we describe the value function, we will provide some intuition for the volatility constraint and describe in more detail the information-cost matrix function. The volatility 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.

Our specification assumes that her beliefs are driven by a Brownian motion, which generates continuous sample paths and does not have jumps.[11] This embeds the idea that, as one looks at smaller and smaller time intervals, the informativeness of the signals the DM is observing scales down. In section 6, we discuss more primitive assumptions about the cost of alternative dynamic information sampling strategies that lead the DM to want to smooth the quantity of information gathered across time, so that the continuity assumed

[10]At this point, we have also not shown that $V(q)$ is differentiable everywhere, but this is proven in the proof of Theorem 1.

[11]Che and Mierendorff (2016) and Zhong (2017) explore related models with jumps in beliefs. These are assumed to represent the only possible form of information arrival in the former paper, and demonstrated to represent an optimal form of experimentation in the latter paper, under assumptions different from those made here.

in this section is a feature of the optimal strategy, in a continuous-time limiting case of the model presented in that section.

We derive the information constraint (equation (2)) from a model in which the DM can choose any information structure she desires at each time period, as in standard rational inattention models. One result of our derivation is the observation 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.

In modeling the evolution of the DM's beliefs as a diffusion process, our model resembles those proposed by authors such as Krajbich et al. (2014) and Fudenberg et al. (2015), though 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 costly to obtain more precise information of some kinds rather than others. Larger (positive) diagonal elements k_{xx} for certain states x imply that it is relatively more costly to obtain signals that reveal much about the likelihood of those states; larger negative off-diagonal elements $k_{xx'}$ (relative to the size of the diagonal elements k_{xx} and $k_{x'x'}$) for pairs of states x, x' imply that it is relatively more costly to obtain signals that allow one to differentiate sharply between states x and x'.

An example of an information-cost matrix function that satisfies our general assumptions (and will be important for the discussion below) is the inverse Fisher information matrix $(g^{+}(q))$,[12]

$$k(q) = g^{+}(q) = \begin{bmatrix} q_1(1-q_1) & -q_1 q_2 & \cdots & -q_1 q_{|X|} \\ -q_1 q_2 & q_2(1-q_2) & \cdots & -q_2 q_{|X|} \\ \vdots & \vdots & \ddots & \vdots \\ -q_1 q_{|X|} & -q_2 q_{|X|} & \cdots & q_{|X|}(1-q_{|X|}) \end{bmatrix}. \tag{4}$$

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. Thus any two states are assumed to be equally easy or difficult to tell apart: it only matters whether two states are the same or not, and how likely they are to occur.

While this kind of symmetry might seem appealing on a priori grounds for some applications (where the different possible states are a set of alternatives, each equally unrelated to all of the others), we view it as quite implausible for many cases of economic relevance. For example, one is often interested in states that represent different possible values of some quantity (a "state variable"), and hence can be ordered on a line.[13] One might well suppose that possible methods of learning about the value of that variable will all have the

[12]The Fisher information matrix, of which this can be viewed as a pseudo-inverse, is described in section 4.2.

[13]This is also often true of perceptual classification experiments, in which subjects are asked to classify stimuli that differ from one another in their intensity or magnitude along some single dimension.

property that nearby values of the state variable result in similar probabilities of receiving particular signals, and hence that it is particularly costly to arrange an information structure that makes the conditional probabilities of signals very different for states that are near one another in the ordering of states.

An alternative possible information-cost matrix function, also satisfying our general assumptions, is given by

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

In this case, the only off-diagonal elements $k_{xx'}(q)$ are negative elements in the case that 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 conditional probabilities of signals are different ($p_{x'} \neq p_x$).

The differing implications of these two alternative assumptions about the form of the information-cost matrix function are explored in section 5. For now, we simply note that our model allows for different specifications in this regard, and that we regard this as desirable, as it will often be reasonable for the specification of information costs to incorporate a notion of "distance" between different possible states.

Our derivation of the continuous-time problem set out above from a more explicit evidence accumulation problem places additional restrictions on the information-cost matrix function, beyond the properties already mentioned above: it will in fact be necessary that $k(q)$ be continuous, and that there exist a positive constant m such that $k(q) - mg^{+}(q)$ is positive semi-definite.[14]

For a large class of information-cost matrix functions $k(q_t)$, we can solve the 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.

3 The Equivalence of 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 to be 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, it is possible to establish an equivalence between the information that is acquired through an optimal evidence accumulation process of the kind proposed in 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 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.

[14]Examples (4) and (5) above are both continuous in q. The second of these examples does not strictly satisfy the second requirement stated in the text for $m > 0$, but is the limit of a sequence of examples that does. These examples are closely related to the mutual-information cost function proposed by Sims and to a "neighborhood-based" cost function that we introduce in section 5, respectively.

3.1 Static Rational Inattention Problems

We begin by explaining the form of a static rational inattention problem. As in the previous section, let $x \in X$ be the underlying state of nature, and let $s \in S$ be a signal the DM can receive, which might convey information about the state. We assume that X and S are finite sets. Let $q \in \mathscr{P}(X)$ denote the DM's prior belief (before receiving a signal) about the probability of state x. Define $p_{s,x}$ as the probability of receiving signal s in state x, let $p_x \in \mathscr{P}(S)$ be the associated conditional probability distribution of the signals given state x, and let p be the $|S| \times |X|$ matrix whose elements are $p_{s,x}$. The matrix p, which is a set of conditional probability distributions for each state of nature, $\{p_x\}_{x \in X}$, defines an "information structure." After receiving signal s, the DM will hold a posterior, $q_s \in \mathscr{P}(X)$, which is a function of p and q, defined by Bayes' rule.

The maximum achievable expected payoff, given an information structure p and prior q, can be written as

$$\bar{u}(p,q) \equiv \max_{\{a(s)\}} \sum_{x \in X} \sum_{s \in S} q_x p_{s,x} u(a(s),x).$$

The standard static rational inattention problem, given the signal alphabet S,[15] is then

$$\max_{\{p_x \in \mathscr{P}(S)\}_{x \in X}} \bar{u}(p,q) - \theta C(p,q;S), \tag{6}$$

where

$$C(\cdot,\cdot;S) : \mathscr{P}(S)^{|X|} \times \mathscr{P}(X) \to \mathbb{R} \tag{7}$$

is a cost function for information structures, and $\theta > 0$ is a multiplicative factor that lets us consider alternative assumptions about the tightness of the information constraint, given a measure of the informativeness of alternative information structures represented by the

[15] The 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.

function C.

In the classic formulation of Sims, a problem of the form equation (6) is considered, in which the cost function $C(p,q;S)$ is given by the Shannon mutual information between the signal and the state. This can be defined using Shannon's entropy,[16]

$$H^{Shannon}(q) \equiv -\sum_{x \in X} (e_x^T q) \ln(e_x^T q). \tag{8}$$

Shannon's entropy can in turn be used to define a measure of the degree to which the posterior q_s associated with any signal differs from the prior q, the Kullback-Leibler (KL) divergence,

$$D_{KL}(q_s||q) \equiv H^{Shannon}(q) - H^{Shannon}(q_s) + (q_s - q)^T H_q^{Shannon}(q). \tag{9}$$

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

$$I^{Shannon}(p,q;S) \equiv \sum_{s \in S} (e_s^T pq) D_{KL}(q_s||q). \tag{10}$$

It is a measure of the informativeness of the signal, in that it provides a measure of the degree to which the signal changes what one should believe about the state, on average.

Shannon's 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 discuss additional examples below.

[16]We use the notation e_x to denote the vector (element of $\mathbb{R}^X$) with a one in the place corresponding to state x, and zeros elsewhere (column x of the identity matrix of dimension $|X|$).

3.2 A Tractable Class of Continuous-Time Models

We return to our discussion of the continuous-time information sampling problem introduced in section 2. To obtain further results, we restrict our attention to information-cost matrix functions with the following property: there exists a twice-differentiable function $H : \mathbb{R}_+^{|X|} \to \mathbb{R}$ such that, for all q_t in the interior of the simplex,

$$D(q_t)^{-1} k(q_t) D(q_t)^{-1} = H_{qq}(q_t). \tag{11}$$

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 5.[17] We shall refer to the function H as the "generalized entropy function," for reasons that will become clear below. Using this convex function, we can define a Bregman divergence,

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

The Kullback-Leibler divergence is a Bregman divergence (see equation (9)), with a generalized entropy function equal to the negative of Shannon's entropy.

With these information-cost matrix functions, it is easy to show (using equation (3)) that the quantity $V(q) - \theta H(q)$ is a local martingale inside the continuation region, anywhere the value function is twice differentiable. Ignoring several technicalities, which are

[17] It is more restrictive, however, than the class of information-cost matrix functions defined in section 2.

discussed in the proof, we can apply the optional stopping theorem:

$$\begin{aligned} V(q_0) &= E_0[V(q_\tau) - \theta H(q_\tau) + \theta H(q_0)] \\ &= E_0[\hat{u}(q_\tau) - \theta H(q_\tau) + \theta H(q_0)]. \end{aligned}$$

Using this idea, and the notion that, in an optimal stopping problem, the DM "chooses the boundaries," we conjecture and verify the following result:

Theorem 1. *There exists a unique solution to the continuous time sequential evidence accumulation problem, in which*

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

subject to the constraint that $\sum_{a \in A} \pi(a) q_a = q_0$.

There exist maximizers of this problem, π^* *and* q_a^**, such that* π^* *is the unconditional probability, in the dynamic 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 A.2. □

Thus the sequential evidence accumulation problem is equivalent to a static rational inattention problem of the kind stated in section 3.1, with a particular kind of static information-cost function,

$$\begin{aligned} C(p, q_0; S) &= \sum_{s \in S} \pi(s) D_H(q_s || q_0), \\ &= \sum_{s \in S} \pi(s) H(q_s) - H(q_0) \end{aligned} \tag{12}$$

where $\pi(s)$ now refers to the unconditional probability of receiving signal s in the static problem, and q_s to the posterior when signal s is received; and with the signal space S in the static problem identified with the set of possible actions A.[18] We shall call a cost function that can be written in the form (12) "posterior-separable."[19]

The mutual-information cost function (10) proposed by Sims is one such cost function. In this case, the generalized entropy function H is the negative of Shannon's entropy (8), the corresponding information-cost matrix function is the inverse Fisher information matrix (4), the Bregman divergence is the Kullback-Leibler divergence (9), and the information measure defined by (12) is then the Shannon mutual information (10). Thus Theorem 1 provides a foundation for assuming endogenous information of the same kind as the standard static rational inattention model, and hence for the same predictions regarding stochastic choice as are obtained by Matêjka et al. (2015).[20] This assumes a particular information-cost matrix function; but below we show not only that a continuous-time model with this particular information-cost matrix function can arise as the limit of a well-behaved discrete-time model of sequential evidence accumulation, but also that any of an entire class of possible specifications of the flow information cost function for that discrete-time model will lead to this result in the continuous-time limit.

On the other hand, Theorem 1 also implies that other posterior-separable cost functions can similarly be justified. Indeed, any static information cost function (12), where D_H is the Bregman divergence derived from some convex, twice-differentiable function H, can be given such a justification.[21] We give additional examples in section 5. The divergence D_H associated with such a model is of interest apart from its role in defining the equivalent

[18]The "signal" can thus be viewed as an instruction as to which action is advisable.

[19]This kind of cost function is instead called "uniformly posterior-separable" in Caplin et al. (2017). The class of static cost functions that can be justified by Theorem 1 is also related to the class of "GERI" cost functions defined by Fosgerau et al. (2016).

[20]For a related foundation for this static cost function, in the special case in which there are only two possible states, see Morris and Strack (2017).

[21]The continuous-time information sampling process that is required is simply the one in which the information-cost matrix function is given by equation (11).

static information-cost function (12). In particular, the expected value of the divergence indicates the expected time cost required for the DM to reach a decision.

Theorem 1 shows that we can associate a particular posterior-separable static information-cost function with any matrix-valued function $k(q)$ satisfying certain conditions: this is the information cost function that defines a static rational-inattention problem that is equivalent in certain respects to the continuous-time problem defined by $k(q)$. But we also show below how to derive a particular information-cost matrix function $k(q)$ corresponding to any information-cost function $C(p,q;S)$ satisfying certain conditions: $k(q)$ defines a local approximation to the function $C(p,q;S)$, and defines a continuous-time problem that represents a limiting case of a discrete-time optimal evidence accumulation problem, in which $C(p,q;S)$ is the flow information-cost function specifying the cost associated with each individual signal that is received.

Thus, there is a two-way relationship between matrix-valued functions $k(q)$ and cost functions $C(p,q;S)$. Posterior-separable cost functions are the fixed points of the resulting mapping: if one uses a posterior-separable cost function to derive an information-cost matrix function, and then solves the continuous-time model defined by that function $k(q)$, one will recover that same posterior-separable cost function as the static information-cost function of the equivalent static rational inattention problem.

Another important general consequence of Theorem 1 concerns posterior beliefs at the time of an eventual decision. The probability distribution $q_a^* \in \mathscr{P}(X)$ is the DM's belief conditional on taking action $a \in A$. The vector q_a^* is unique, given a particular action a, meaning that there is only one belief the DM can reach before choosing to stop and take a particular action. (The further this belief is from the DM's prior, q_0, as measured by the divergence D_H, the more time it will take, in expectation, for the DM to arrive at this belief before acting.) The martingale property of beliefs during the evidence accumulation process thus requires that beliefs q_t at each stage of the process are some convex combination of the finite set of posteriors $\{q^*_a\}_{a\in A}$.

Hence beliefs diffuse on a simplex of dimension $|A|-1$ during the decision process; if there are only two possible actions (as in the binary choice problems to which the drift-diffusion model is applied by authors such as Fehr and Rangel (2011)), then the belief state must diffuse along a line segment, as assumed in the DDM, regardless of the number of possible states $|S|$. Thus an apparently arbitrary feature of the DDM (outside the two-state case in which the DDM is known to correspond to optimal Bayesian decision making, as discussed by Fudenberg et al. (2015)) can be shown to be follow from optimal sequential evidence accumulation, if the information sampling is flexible in the way that we model it here.

4 Flow Information Costs

In this section, we further elaborate on the connection, suggested above, between the information-cost matrix function in our continuous-time model of information sampling and the kind of cost function for an individual signal that is assumed in static rational inattention models. The continuous-time model can be viewed as a limiting case of a discrete-time model of optimal evidence accumulation, in which an endogenously determined signal is received each period, and the choice of the signal to receive each period is subject to a cost for more informative signals, specified by an information-cost function similar to the kind assumed in static rational inattention models. We call this cost function for an individual signal in a dynamic model of evidence accumulation the "flow information-cost function."

While we defer until section 6 a complete discussion of how the continuous-time model can be derived as a limiting case of a discrete-time dynamic model, in this section we preview certain conclusions from that analysis by explaining the connection between the flow information-cost function of the discrete-time dynamic model and the information-cost matrix function of the continuous-time model. (Here the connection is simply asserted; the connection is proven in section 6.) We preview the results before proceeding to the complete derivation, because understanding them can help to explain the assumptions about the information-cost matrix function that we have proposed above. Additionally, in the following section we discuss a particular class of information-cost matrix functions, and wish to motivate this specification in terms of the form of flow information-cost functions from which these information-cost matrix functions can be derived.

An important conclusion of this section is our demonstration that many different flow information-cost functions can give rise to the same information-cost matrix function, and hence to the same predictions about the information that will be accumulated in the continuous-time limit. This is one of the main advantages, in our view, of considering the continuous-time limit: while our conclusions still depend on assumptions about the nature of information costs, there are less ways in which our conclusions can vary once we pass to the continuous-time limit.

4.1 Assumptions about Flow Information-Cost Functions

At each stage of the discrete-time sequential evidence accumulation problem discussed in section 6, the DM chooses an information structure that determines the kind of signal observed in that stage. Any such information structure p has a cost $C(p,q;S)$, given by a function of the form of (7), where q indicates the DM's prior in this stage (that is, the posterior beliefs following from observations prior to the current stage of the dynamic problem), and S is again the signal alphabet.[22] Our most general results depend only on

[22]The information-cost functions that we study, like mutual information, are defined for all finite signal alphabets S. Note, however, that mutual information is also defined over alternative sets of states of nature X. We do not impose this requirement on our more general cost functions — all of our analysis takes the set of states of nature as given.

assuming that this flow information-cost function satisfies a set of six general conditions, stated below.

All of these conditions are satisfied by the mutual-information cost function (10) proposed by Sims, but they are also satisfied by many other cost functions. (Additional examples are given in section 5.) They are closely related to conditions that other authors have also proposed as attractive general properties to assume about information-cost function, though in the context of static information-cost functions of the kind discussed in section 3.1. Here we *assume* that the flow information-cost function in our dynamic model satisfies all six of these conditions; we then *prove* that under our assumptions, the equivalent static rational inattention problem (the existence of which is guaranteed by Theorem 1) involves a static information-cost function that satisfies these conditions.

Condition 1. Information structures that convey no information ($p_x = p_{x'}$ for all x, x' in the support of q) have zero cost. All other information structures have a weakly positive cost.

This condition ensures that the least costly strategy for the DM in the standard static rational inattention problem is to acquire no information, and make her decision based on the prior. The requirement that gathering no information has zero utility cost is a normalization.

The next condition is called mixture feasibility by Caplin and Dean (2015). Consider two information structures, $\{p_{1,x}\}_{x \in X}$, with signal alphabet S_1, and $\{p_{2,x}\}_{x \in X}$, with alphabet S_2. Given a parameter $\lambda \in (0,1)$, we define a mixed information structure, $\{p_{M,x}\}_{x \in X}$ over the signal alphabet $S_M = (S_1 \cup S_2) \times \{1,2\}$. For each $s = (s_1, 1)$ in the alphabet S_M, $p_{M,x}(s)$ is equal to $\lambda p_{1,x}(s)$ if $s_1 \in S_1$, and equal to 0 otherwise. Likewise, for each $s = (s_2, 2)$, $p_{M,x}(s)$ is equal to $(1-\lambda) p_{2,x}(s)$ if $s_2 \in S_2$, and equal to 0 otherwise.

That is, this information structure results, with probability λ, in a posterior associated with information structure p_1, and with probability $1-\lambda$ in a posterior associated with information structure p_2. The distribution of posteriors under the mixed information structure is a convex combination of the distributions of posteriors under the two original information structures, as if the DM flipped a coin, observed the result, and then randomly chose one of the two information structures. The mixture feasibility condition requires that choosing a mixed information structure costs no more than the cost of randomizing over information structures (using a mixed strategy in the rational inattention problem).

Condition 2. Given two information structures, $\{p_{1,x}\}_{x\in X}$, with signal alphabet S_1, and $\{p_{2,x}\}_{x\in X}$, with alphabet S_2, the cost of the mixed information structure is weakly less than the weighted average of the cost of the separate information structures:

$$C(p_M, q; S_M) \leq \lambda C(p_1, q; S_1) + (1-\lambda)C(p_2, q; S_2).$$

The next condition uses Blackwell's ordering. Consider two signal structures, $\{p_x\}_{x\in X}$, with signal alphabet S, and $\{p'_x\}_{x\in X}$, with alphabet S'. The first information structure Blackwell dominates the second information structure if, for all utility functions $u(a,x)$ and all priors $q \in \mathscr{P}(X)$,

$$\bar{u}(p,q) \;\geq\; \bar{u}(p',q).$$

If one information structure Blackwell dominates another, it is weakly more useful for every decision maker, regardless of that decision maker's utility function and prior. In this sense, it conveys weakly more information. This ordering is incomplete; most information structures neither dominate nor are dominated by a given alternative information structure. However, when an information structure does Blackwell dominate another one, we assume that the dominant information structure is weakly more costly.

Condition 3. If the information structure $\{p_x\}_{x\in X}$ with signal alphabet S is more informative, in the Blackwell sense, than $\{p'_x\}_{x \in X}$, with signal alphabet S', then, for all $q \in \mathscr{P}(X)$,

$$C(\{p_x\}_{x \in X}, q; S) \geq C(\{p'_x\}_{x \in X}, q; S').$$

The first three conditions are, from a certain perspective, almost innocuous. For any joint distribution of actions and states that could have been generated by a DM solving a rational inattention type problem, with an arbitrary information cost function, there is a cost function consistent with these three conditions that also could have generated that data (Theorem 2 of Caplin and Dean (2015)). The result arises from the possibility of the DM pursuing mixed strategies over information structures, or in the mapping between signals and actions. These conditions also characterize "canonical" rational inattention cost functions, in the terminology of De Oliveira et al. (2017).

The mixture feasibility condition (Condition 2) and Blackwell monotonicity condition (Condition 3) are equivalent to requiring that the cost function be convex over information structures and Blackwell monotone. We summarize this equivalence in the following lemma.

Lemma 2. *Let p and p' be information structures with signal alphabet S. A cost function is convex in information structures if, for all $\lambda \in (0,1)$, all signal alphabets S, and all $q \in \mathscr{P}(X)$,*

$$C(\lambda p + (1-\lambda)p', q; S) \leq \lambda C(p, q; S) + (1-\lambda) C(p', q; S).$$

A cost function satisfies mixture feasibility and Blackwell monotonicity (Conditions 2 and 3) if and only if it is convex in information structures and satisfies Blackwell monotonicity.

Proof. See the appendix, section A.3. □

The fourth condition that we assume, which is not imposed by Caplin and Dean (2015), Caplin et al. (2017), or De Oliveira et al. (2017), is a differentiability condition that will allow us to characterize the local properties of our cost functions.

Condition 4. For all signal alphabets S, in a neighborhood around any uninformative information structure, the information cost function is continuously twice-differentiable in information structures $\{p_x\}_{x\in X}$, in all directions that do not change the support of the signal distribution, and directionally differentiable, with continuous directional derivatives, with respect to perturbations that increase the support of the signal distribution. The information cost function is also Lipschitz-continuous in q.

While this may seem a relatively innocuous regularity condition, it is not completely general; for example, it rules out the case in which the DM is constrained to use only signals in a parametric family of probability distributions, and the cost of other information structures is infinite. Thus it rules out information structures of the kind assumed in Fudenberg et al. (2015) or Morris and Strack (2017). Condition 4 also rules out other proposed alternatives, such as the channel-capacity constraint suggested by Woodford (2012).

The next condition that we assume, which is also not imposed by Caplin and Dean (2015), Caplin et al. (2017), or De Oliveira et al. (2017), is a sort of local strong convexity. We will assume that the cost function exhibits strong convexity, in the neighborhood of an uninformative information structure, with respect to information structures that hold fixed the unconditional distribution of signals, uniformly over the set of possible priors.

Condition 5. There exists constants $m > 0$ and $B > 0$ such that, for all priors $q \in \mathscr{P}(X)$, and all information structures that are sufficiently close to uninformative ($C(p,q;S) < B$),

$$C(p,q;S) \geq \frac{m}{2} \sum_{S\in S} (e_s^T pq)||q_s - q||_X^2,$$

where q_s is the posterior given by Bayes' rule and $||\cdot||_X$ is an arbitrary norm on the tangent space of $\mathscr{P}(X)$.

This condition is slightly stronger than Condition 1; it it essentially an assumption of "local strong convexity" instead of merely local convexity. It implies that all informative information structures have a strictly positive cost, and that (regardless of the DMs' current beliefs) there are no informative information structures that are "almost free."

The mutual-information cost function (10) satisfies each of these five conditions. However, it is not the only cost function to do so. For example, we can construct a family of such cost functions, using the family of "f-divergences," defined as

$$D_f(q_s||q) = \sum_{x \in X} (e_x^T q) f(\frac{e_x^T q_s}{e_x^T q}),$$

where f is any strictly convex, twice-differentiable function with $f(1) = f'(1) = 0$ and $f''(1) = 1$. (The KL divergence is a member of this family, corresponding to $f(u) = u \ln u - u + 1$.) For any divergence in this family, we can define an information cost function

$$I_f(p,q;S) = \sum_{s \in S} (e_s^T pq) D_f(q_s||q). \tag{13}$$

(When I_f is the KL divergence, this is just mutual information.) It is relatively easy to observe that this family of information cost functions satisfies all five of the conditions described above.[23]

As another example of a class of cost functions that satisfy the conditions, we can establish the following.

Corollary 1. *Under the assumptions of Theorem 1, the posterior-separable cost function (12) that defines the equivalent static rational inattention problem satisfies Conditions 1-5.*

[23]This follows from Lemma 3 in the next section.

This follows directly from the form of the cost function (12) and Lemma 3, described in the next section.

4.2 Local Characterization of Flow Information Costs

Next we discuss the local (second-order) properties of any information cost function satisfying the conditions stated above. The condition requiring that Blackwell-dominant information structures cost weakly more (Condition 3) is of particular importance. To understand why, it is first useful to recall Blackwell's theorem.

Theorem. *(Blackwell (1953)) The information structure $\{p_x\}_{x\in X}$, with signal alphabet S, is more informative, in the Blackwell sense, than $\{p'_x\}_{x\in X}$, with signal alphabet S', if and only if there exists a Markov transition matrix $\Pi : S \to S'$ such that, for all $s' \in S'$ and $x \in X$,*

$$p'_x = \Pi p_x. \tag{14}$$

This Markov transition matrix is known as the "garbling" matrix. Another way of interpreting Condition 3 is that garbled signals are (weakly) less costly than the original signal.

There are certain kinds of garbling matrices that don't really garble the signals. These garbling matrices have left inverses that are also Markov transition matrices. If we define an information structure $\{p_x\}_{x\in X}$, with signal alphabet S, and another information structure $\{p'_x\}_{x\in X}$, with signal alphabet S', using one of these left-invertible matrices, via equation (14), then $\{p_x\}_{x\in X}$ is more informative than $\{p'_x\}_{x\in X}$, but $\{p'_x\}_{x\in X}$ is also more informative than $\{p_x\}_{x\in X}$. These two information structures are called "Blackwell-equivalent," and it follows that the cost of these two information structures must be equal, by Condition 3. The left-invertible Markov transition matrices associated with Blackwell-equivalent information structures are called Markov congruent embeddings by Chentsov

(1982). Chentsov (1982) studied tensors and divergences that are invariant to Markov congruent embeddings (we will say "invariant" for brevity).

An invariant divergence is a divergence that is invariant to these embeddings. Let Π be a Markov congruent embedding from $\mathscr{P}(S)$ to $\mathscr{P}(S')$. The KL divergence and the f-divergences more generally are invariant, meaning that

$$D_f(\Pi p||\Pi r) = D_f(p||r)$$

for all $p,\, r \in \mathscr{P}(S)$. There are also other, non-additively-separable invariant divergences. Chentsov's theorem (Chentsov (1982)) states that, for any invariant divergence D_I,

$$\frac{\partial^2 D_I(p||r)}{\partial p^i \partial p^j}|_{p=r} = c \cdot g_{ij}(r), \tag{15}$$

where $c > 0$ is a positive constant and $g_{ij}(r)$ is the (i,j)-element of the Fisher information matrix evaluated at r. Written in terms of the coordinate system used previously in the paper,[24]

$$g(r) = D(r)^{+} - \iota\iota^T. \tag{16}$$

Here, $D^{+}(r)$ denotes the inverse of the $|S| \times |S|$ diagonal matrix whose diagonal elements are the elements of r, and ι is a vector of ones of length $|S|$.

However, the focus of this paper is not invariant divergences, but rather invariant information cost functions. By Condition 3, all information cost functions satisfying our conditions are invariant to Markov congruent embeddings. It necessarily follows that, for any Markov congruent embedding Π, that

[24] This corresponds to the standard definition of the Fisher information matrix, if derivatives of smooth functions defined on the probability simplex are written in terms of the coordinates explained in footnote 7.

$$C(\{p_x\}_{x\in X}, q; S) = C(\{\Pi p_x\}_{x\in X}, q; S').$$

Using this invariance, and results from Chentsov (1982), we will describe the local structure of all information cost functions satisfying our conditions.

Chentsov establishes the following results:[25]

i) Any continuous function that is invariant over the probability simplex is equal to a constant.

ii) Any continuous, invariant 1-form tensor field over the probability simplex is equal to zero.

iii) Any continuous, invariant quadratic form tensor field over the probability simplex is proportional to the Fisher information matrix.[26]

These results allow us to characterize the local properties of rational inattention cost functions, via a Taylor expansion. Hold fixed the signal alphabet S, and consider an information structure $p_x(\varepsilon, \nu) = r + \varepsilon\tau_x + \nu\omega_x$, where $r \in \mathscr{P}(S)$. Here, τ_x satisfies $\iota^T\tau_x = 0$ for all x, and. for all $s \in S$, $e_s^T\tau_x \neq 0$ only if $e_s^T r > 0$. That is, τ_x is an element of the tangent space of the probability simplex at r, and the same holds true for ω_x. As a result, for values of the perturbation parameters ε and ν sufficiently close to zero, $p_x \in \mathscr{P}(S)$ for all $x \in X$. In other words, the parameters ε and ν index a two-parameter family of perturbations of an uninformative information structure (corresponding to $\varepsilon = \nu = 0$), in which the perturbed

[25]See Lemma 11.1, Lemma 11.2, and Theorem 11.1 in Chentsov (1982). See also Proposition 3.19 of Ay et al. (2014), who demonstrate how to extend the Chentsov results to infinite sets X and S.

[26]A 1-form tensor field on a probability simplex $\mathscr{P}$ is a function $T : V \times \mathscr{P} \to \mathbb{R}$, where V is the tangent space of the simplex. Let $\Pi : \mathscr{P} \to \mathscr{P}'$ be a mapping from the simplex $\mathscr{P}$ to the simplex $\mathscr{P}'$, let V' be the tangent space of the simplex $\mathscr{P}'$, and let $d\Pi : V \to V'$ be the pushforward of the mapping Π. The tensor field is invariant under Π if $T(d\Pi v, \Pi p) = T(v, p)$ for all $p \in \mathscr{P}$ and v in the tangent space at p, and a similar definition holds for quadratic form tensor fields. This means that if the quadratic form is extended to a quadratic form over $\mathbb{R}^{|X|}$ using the method defined in footnote 6, its matrix representation is proportional to the matrix defined in (16).

information structures will generally be informative; the τ_x and ω_x specify two directions of perturbation. Each of the perturbed information structures has the property that p_x is absolutely continuous with respect to r.

By Condition 1, $C(\{p_x(0,0)\}_{x\in X};q;S) = 0$. The first order term is

$$\frac{\partial}{\partial \varepsilon} C(\{p_x(\varepsilon,\nu)\}_{x\in X}, q; S)|_{\varepsilon=\nu=0} = \sum_{x\in X} C_x(\{r\}_{x\in X}, q; S) \cdot \tau_x,$$

where C_x denotes the derivative with respect to p_x. This derivative, $C_x(\{r\};q;S)$, forms a continuous 1-form tensor field over the probability simplex $\mathscr{P}(S)$. By the invariance of $C(\cdot)$, it also follows that C_x is invariant, and therefore, by Chentsov's results, it is equal to zero.

We repeat the argument for the second derivative terms. Those terms can be written as

$$\frac{\partial}{\partial \nu}\frac{\partial}{\partial \varepsilon} C(\{p_x(\varepsilon,\nu)\}_{x\in X}, q; S)|_{\varepsilon=\nu=0} = \sum_{x'\in X}\sum_{x\in X} \omega_{x'}^T \cdot C_{xx'}(\{r\}_{x\in X}, q; S) \cdot \tau_x.$$

By the invariance of $C(\cdot)$, the quadratic form $C_{xx'}(\cdot)$ is invariant for all $x,x' \in X$, and therefore is proportional to the Fisher information matrix for all $x,x' \in X$. We can define a matrix $k(q)$ consisting of the constants of proportionality associated with each $x,x' \in X$.[27] That is,

$$\frac{\partial}{\partial \nu}\frac{\partial}{\partial \varepsilon} C(\{p(\cdot|\cdot;\varepsilon,\nu)\}, q)|_{\varepsilon=\nu=0} = \sum_{x'\in X}\sum_{x\in X} (e_x^T k(q) e_{x'}) \omega_{x'}^T g(r) \tau_x,$$

where $g(r)$ is the Fisher information matrix evaluated at the unconditional distribution of signals $r \in \mathscr{P}(S)$. We note that the matrix $k(q)$ can depend on the prior q, but cannot depend on the unconditional distribution of signals, r; otherwise, invariance would not hold.

In the case of the mutual-information cost function, the matrix $k(q)$ is itself the inverse

[27] Our re-use of the notation $k(q)$ here is intentional.

Fisher information matrix,

$$k(q) = g^{+}(q) = D(q) - qq^T.$$

In general, however, the matrix-valued function $k(q)$ is not the inverse Fisher information matrix, but rather an arbitrary matrix-valued function satisfying certain restrictions. Our choice of notation hints at what we will prove in section 6: that this matrix-valued function is just the information-cost matrix function described in section 2.

We are now in position to discuss our approximation of the information cost function. We use Taylor's theorem to approximate the cost function and its gradient up to order Δ (we use Δ because in section 6, we will be looking at small time intervals). We consider perturbations that, as above, preserve the support of the signal structure. As a result, this theorem should be interpreted as applying to "frequent but not very informative" signals, as opposed to "rare but informative" signals. We will discuss the latter type of signals shortly.

Theorem 2. *Suppose that an information structure $\{p_x\}_{x\in X}$, with signal alphabet S, is described by the equation*

$$p_x = r + \Delta^{\frac{1}{2}}\tau_x + o(\Delta^{\frac{1}{2}}),$$

where, for any $x \in X$ and any $\Delta \geq 0$, $e_s^T p_x \neq 0 \Rightarrow e_s^T r > 0$. Let $C(\cdot)$ be an information cost function that satisfies Conditions 1-4. Then, for Δ sufficiently small, the cost of this information structure can be written as

$$C(\{p_x\}_{x\in X}; q; S) = \frac{1}{2}\Delta \sum_{x'\in X}\sum_{x\in X}(e_x^T k(q) e_{x'})\tau_{x'}^T g(r)\tau_x + o(\Delta),$$

where the matrix $k(q)$ is positive semi-definite and symmetric, and satisfies $k(q)\iota = 0$.

If in addition the cost function satisfies Condition 5, Then there exists a constant $m_g > 0$ such that the difference between $k(q)$ and the inverse Fisher information matrix, $g^{+}(q)$, *multiplied by that constant, is positive semi-definite:* $k(q) - m_g g^+(q) \succeq \mathbf{0}$.

Proof. See appendix, section A.4. □

There are, in effect, two ways for a signal to be contain a small amount of information, and different costs associated with these different types of signals. The results of Theorem 2 characterize, for any rational inattention cost function satisfying our conditions, the cost of receiving frequently, but relatively uninformative, signals. As Corollary 2 below demonstrates, the posteriors associated with these signals are close to the prior (order $\Delta^{\frac{1}{2}}$). We will discuss the cost of receiving a rare but informative signal below. Previewing the results of section 6, these two types of uninformative signals correspond, in the continuous-time limit, to the diffusion and jump components of the belief process.

The theorem substantially restricts the local structure of the cost of commonly occurring, but not particularly informative, signals, relative to the most general possible alternatives (which would not satisfy our conditions). Potential information structures $\{p_x\}_{x \in X}$ can be represented as vectors of dimension $N = (|S| - 1) \times |X|$. Under the assumptions of Condition 1, convexity, and Condition 4 (but not the Blackwell ordering condition, Condition 3), the cost function must locally resemble an inner product with respect to a positive semi-definite, $N \times N$ matrix. If we impose Condition 3 as well, the results of Theorem 2 show that we can restrict this matrix to the $k(q)$ matrix, an $|X| \times |X|$ matrix. If the DM were only allowed binary signals ($|S| = 2$), this restriction would be trivial. When the DM is allowed to contemplate more general information structures, the restriction is non-trivial.

Several authors (Caplin and Dean (2015); Caplin et al. (2017); Kamenica and Gentzkow (2011)) have observed that it is easier to study rational inattention problems by considering the space of posteriors, conditional on receiving each signal, rather than space of signals. We can redefine the cost function using the posteriors and unconditional signal probabilities, rather than the prior and the conditional probabilities of signals. The results are described in the corollary below.

Corollary 2. *Under the assumptions of Theorem 2, the posterior beliefs can be written, for any $s \in S$ such that $e_s^T r > 0$, as*

$$q_{s,n,x} = q_x + \Delta^{\frac{1}{2}} q_x \frac{e_s^T \tau_{n,x}}{e_s^T r} + o(\Delta^{\frac{1}{2}}).$$

Define the matrix

$$\bar{k}(q) = D^+(q) k(q) D^+(q),$$

where $D^+(\cdot)$ is the pseudo-inverse of the diagonal matrix. The cost function can be written as

$$C(\{p_x\}_{x\in X}, q; S) = \frac{1}{2} \sum_{s\in S: e_s^T r>0} (e_s^T r)(q_s - q)^T \bar{k}(q)(q_s - q) + o(\Delta).$$

Proof. See the appendix, section A.5. □

The theorem and corollary above describe the costs of receiving frequent, but relatively uninformative signals. We next discuss the cost of receiving rare, but informative signals. These types of signals, in the limit that we discuss in section 6, will lead to jumps in beliefs. After we describe the cost of these signals, we will introduce a condition that ensures that jumps in beliefs are not part of an optimal evidence accumulation strategy.

Corollary 3. *Under the assumptions of Theorem 2, define the signal structure*

$$\hat{p} = \bar{p}_\Delta + \Delta \omega,$$

where p_Δ is a signal structure of the type described in Theorem 2, with $\lim_{\Delta\to 0^+} \bar{p}_\Delta = r\iota^T$, and $\sum_{s\in S} \omega e_x = 0$ for all $x \in X$, with $e_s^T \omega e_x \geq 0$ for all $s \in S$ such that $e_s^T \bar{p}_\Delta = 0$.

The cost of this information structure can be written in the form

$$C(p_n;q;S) = \frac{1}{2}\Delta_n \sum_{s\in S: e_s^T r>0} (e_s^T r)(q_s - q)^T \bar{k}(q)(q_s - q) + \sum_{s\in S: e_s^T r=0} (e_s^T \phi) D^*(q_s||q) + o(\Delta),$$

where the divergence D^* is finite and twice-differentiable in its first argument for q' sufficiently close to q, with

$$\frac{\partial^2 D^*(r||q)}{\partial r^i \partial r^j}|_{r=q} = \bar{k}(q). \tag{17}$$

Proof. See the appendix, section A.6. □

The divergence D^* represents the cost of acquiring an infrequent, but potentially informative, signal. Naturally, if the signal is in fact not very informative, this cost must be closely related to the costs of other uninformative signals, which gives rise to the condition on the Hessian of the divergence. Note that the corollary requires that the cost is additive with respect to the other signals being received (at least up to order Δ). The result follows from the directional differentiability of the cost function with respect to signals that occur with zero probability, the continuity of that directional derivative, and invariance.

We now introduce the last condition we will impose on our cost functions. This condition, which is expressed in terms of the $\bar{k}(q)$ matrix-valued function and the divergence D^*, is a sufficient condition to ensure that the discrete-time models that we study in section 6 converge to the model with continuous sample paths (no jumps) described in section 2. The condition reflects an assumption that learning gradually over time, receiving frequent but never very informative signals, is less costly than receiving rare signals that lead to large changes in beliefs when they occur.

Condition 6. The matrix-valued function $\bar{k}(q)$ and divergence D^* associated with the cost function $C(p,q;S)$ satisfy, for all $q,q' \in \mathscr{P}(X)$ with $q' \ll q$,

$$D^*(q'||q) \geq \frac{1}{2}(q'-q)^T(\int_0^1 \bar{k}(sq'+(1-s)q)ds)(q'-q).$$

We will say that a cost function satisfying this condition exhibits a preference for gradual learning. We will call this preference "strict" if the inequality is strict for all $q' \neq q$. If the $\bar{k}(q)$ function is the Hessian of some generalized entropy function (see equation (11)), this condition is equivalent to requiring that

$$D^*(q'||q) \;\geq\; D_H(q'||q), \tag{18}$$

where D_H is the associated Bregman divergence. In the particular case of mutual information, both D^* and D_H are the KL divergence, and the condition is (weakly) satisfied. It is also easy to construct cases in which it is strictly satisfied, as the example below illustrates.[28]

Consider the family of information cost functions built from f-divergences defined in equation (13) above. All of the cost functions in this family resemble mutual information, to second order, in the sense defined by Corollary 2. Assuming that the posteriors induced by the information structure p and prior q, $\{q_s\}_{s\in S}$, are close to the prior q, and that the

[28] It is the assumption that the flow cost function in our dynamic evidence accumulation problem satisfies Condition 6 that allows us to avoid considering the possibility of Poisson jumps in the posterior belief state of the kind assumed by Che and Mierendorff (2016) and Zhong (2017) in the continuous-time model presented in section 2. Zhong (2017) presents conditions under which information accumulation with Poisson jumps can be optimal, but considers only posterior-separable flow cost functions of the form (12) based on a Bregman divergence, so that (18) holds with equality rather than an inequality. In this special case, in our framework jumps can also be among the optimal policies, but an equally good outcome can always be achieved by an information sampling strategy that involves no jumps, as we establish in section 6. When the inequality is instead strict, jumps cannot be optimal in the continuous-time limit of the kind of dynamic evidence accumulation problem considered in this paper.

prior q is on the interior of the simplex,

$$I_f(p,q;S) \approx \frac{1}{2}\sum_{s\in S}(e_s^T pq)(q_s - q)^T D(q)^+ g^+(q) D(q)^+ (q_s - q). \tag{19}$$

In other words, in a sense that we will show formally in section 6, all of these flow cost functions will induce the same information-cost matrix function in the continuous-time problem.

However, all such functions do not induce the same divergence D^*. (Note that for this family, $D^* = D_f$.) Nonetheless, if $D_f \geq D_{KL}$ (which holds strictly, for example, in the case of the χ^2-divergence), these cost functions will generate the same solution in the continuous-time problem: the solution to a static rational-inattention problem with the mutual-information cost function. Regardless of whether the f-divergence used to construct the flow cost function is the KL divergence or not, the KL divergence will appear in the solution to the continuous-time problem. In fact, this result applies to the larger class of invariant divergences, which includes the f-divergences, and follows from Chentsov's theorem (equation (15)).

Of course, we argued in section 2 that the inverse Fisher information matrix, when used as the information-cost matrix function, lacks certain desirable properties related to the distance between different states of the world. In the next section, we will introduce a new family of cost functions, all of which induce information cost matrix functions that do capture these notions. Moreover, these information-cost matrix functions satisfy equation (11), and therefore Theorem 1 applies. We will solve examples of the static model implied by Theorem 1 and compare it to the same static model with mutual information, illustrating why notions of the distance between states matters in economic applications.

5 Neighborhood-Based Cost Functions

Suppose that the state space X can be written as the union of a finite collection of "neighborhoods" $\{X_i\}$, and suppose furthermore that the state space is connected, in the sense that any two states can be connected by a sequence of overlapping neighborhoods. That is, for any two states $x, x' \in X$, there exists a sequence of states $\{x_0, \ldots, x_n\}$ 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. 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 prior $q \in \mathscr{P}(X)$, let $\mathscr{I}(q)$ be the (necessarily non-empty) set of neighborhoods X_i such that some state belonging to X_i has positive probability under the prior, and 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 $q_i \in \mathscr{P}(X_i)$ be the conditional probability distribution over states in neighborhood X_i, given the prior q and conditional on the state being in neighborhood X_i. That is, for all $x \in X_i$,

$$q_i \equiv \frac{1}{\bar{q}_i} E_i q.$$

Similarly, let $q_s \in \mathscr{P}(X)$ be the posterior after receiving signal $s \in S$, and let $q_{i,s} \in \mathscr{P}(X_i)$ be the posterior over states in neighborhood X_i, conditional on receiving signal s and having the state be part of neighborhood X_i. That is, for all $x \in X_i$,

$$q_{i,s} \equiv \frac{1}{\bar{q}_{i,s}} E_i q_s,$$

with $\bar{q}_{i,s} \equiv \sum_{x \in X_i} e_x^T q_s$. We adopt the convention that $q_{i,s} = q_i$ if $\bar{q}_{i,s} = 0$. Finally, let $\bar{p}_i \in \mathscr{P}(S)$ be the conditional distribution of signals under the information structure p and prior q:

$$\bar{p}_i = \frac{\sum_{x \in X_i} p e_x e_x^T q}{\bar{q}_i}.$$

We will say that a cost function has a "neighborhood structure" if it can be written in the form

$$C_N(p,q;S) \;=\; \sum_{i\in\mathscr{I}(q)} \bar{q}_i \sum_{s\in S} e_s^T\,\bar{p}_i\, D_i(q_{i,s}||q_i), \tag{20}$$

where for each $i \in \mathscr{I}(q)$, $D_i(\cdot||\cdot)$ is a divergence (not necessarily the same for all i) defined over probability distributions in $\mathscr{P}(X_i)$ that is a twice-differentiable and strongly convex in its first argument.[29] Mutual information is an example of a flow cost function in this family, corresponding to the case in which there is only a single neighborhood, consisting of the entire state space X, and the divergence is the KL divergence, so that

$$C(p,q;S) \;=\; \sum_{s\in S} (e_s^T pq) D_{KL}(q_s||q) \;=\; I^{Shannon}(\{p_x\},q;S).$$

The information cost functions based on f-divergences, defined by (13), are also single-neighborhood examples of neighborhood-based cost functions.

The following lemma shows that all cost functions with a neighborhood structure satisfy the conditions defined in section 4.1.

Lemma 3. *All cost functions with a neighborhood structure (20) satisfy Conditions 1-4 stated in section 4. If the neighborhood structure includes a neighborhood containing all of the states $x \in X$, the cost function also satisfies Condition 5.*

Proof. See the appendix, section A.7. □

An implication of this lemma is that any posterior-separable cost function (12) based on a strongly convex generalized entropy function H satisfies Conditions 1-5. Below, we give a sufficient condition for Condition 6 to be satisfied as well.

We will study a particular family of cost functions with a neighborhood structure, the "neighborhood-based cost functions." This family is defined by the additional require-

[29] The f-divergences defined previously satisfy these conditions (Amari and Nagaoka (2007)).

ments that (i) the divergences D_i be invariant, and (ii) each of the D_i is bounded below by some positive multiple of D_{KL}, the Kullback-Leibler divergence.[30] As an example of the possibility of satisfying these latter requirements, the D_i may be α-divergences (or Rényi divergences, van Erven and Harremoës (2014))

$$D_\alpha(p_i||q_i) \equiv \frac{1}{\alpha - 1} \log \sum_{x \in X_i} \frac{p_i(x)^\alpha}{q_i(x)^{\alpha-1}},$$

of order $\alpha \geq 1$.[31]

This family can have complex neighborhood structures, for which the requirement that each of the individual divergences D_i be invariant is a less restrictive requirement. The idea of this class of cost functions is that information structures are costly only to the extent that they result in different signal distributions for states that are "similar" to one another, in the sense of belonging to the same neighborhood. If there is only one neighborhood that includes all of the states (the mutual-information case), all states are equally difficult to distinguish from one another. Allowing for more complex neighborhood structures allows us to assume instead that it is much more difficult to tell some pairs of states apart than others. Note that under the general formalism (20), this is true not only because some pairs of states share a neighborhood while others do not — and more generally, that the length of the chain of neighborhoods required to link two states differs for different pairs of states — but also because the divergences D_i can be different for different neighborhoods.

As discussed above, the fact that D_i is an invariant divergence implies that its Hessian

[30] Stipulation (ii) is added in order to ensure that Condition 6 is satisfied. For this it suffices that D_i be bounded below by a Bregman divergence for each i. But as explained in section 4.2, any invariant divergence is locally equivalent (for p near q) to a positive multiple of D_{KL}. Hence in order for D_i to be bounded below by a Bregman divergence, it must be bounded below by a positive multiple of D_{KL}.

[31] The definition is here stated only for the case $\alpha \neq 1$. When $\alpha = 1$, the α-divergence is simply the KL divergence, and Condition 6 is weakly satisfied. If $\alpha > 1$, the α-divergence satisfies $D_\alpha(p||q) > D_{KL}(p||q)$ for all $p \neq q$, so that the strong form of Condition 6 is satisfied, implying a strict preference for gradual learning.

matrix is proportional to the Fisher information matrix. As a result, the approximation described in equation (19) applies, but only within each neighborhood. That is,

$$C_N(p,q;S) \approx \frac{1}{2}\sum_{i\in\mathscr{I}(q)} c_i\bar{q}_i \sum_{s\in S}(e_s^T pq)(q_{i,s}-q_i)^T g(q_i)(q_{i,s}-q_i), \tag{21}$$

where the $c_i > 0$ are positive constants. This implies the following structure for the information-cost matrix:

Lemma 4. *The information-cost matrix function $k_N(q)$ associated with the neighborhood-based cost function is*

$$k_N(q) = \sum_{i\in\mathscr{I}(q)} c_i\bar{q}_i E_i^T g^+(q_i) E_i,$$

where g^+ is the inverse Fisher information matrix and the constant $c_i > 0$ for each i.

Proof. See the appendix, section A.8. □

We can use the information-cost matrix function in our continuous-time problem (the problem defined in section 2).[32] It satisfies the equation necessary for the results of Theorem 1 to apply (equation (11)). As a result, there is a generalized entropy function, $H_N(q)$, associated Bregman divergence, $D_N(p||q)$, and posterior-separable static information-cost function, $C_N^{static}(p,q;S)$, that can be used to define the static rational-inattention problem the choice probabilities of which coincide with the solution to the dynamic model. The following lemma describes these functions:

Lemma 5. *Let $H^{Shannon}(q)$ be Shannon's entropy (8). Then the generalized entropy function $H_N(q)$ associated with the neighborhood-based information-cost matrix function $k_N(q)$*

[32] Our derivation of the continuous-time model from the discrete-time model applies only to cost functions satisfying Conditions 1-6. We have established these conditions only for neighborhood structures that include a neighborhood containing all states. However, the constant c_i associated with this neighborhood can be arbitrarily small, and in what follows we will ignore this requirement.

is given by

$$H_N(q) = -\sum_{i\in\mathscr{I}(q)} c_i \bar{q}_i H^{Shannon}(q_i),$$

and the associated Bregman divergence is

$$D_N(q_s||q) = \sum_{i\in\mathscr{I}(q)} c_i \bar{q}_{i,s} D_{KL}(q_{i,s}||q_i).$$

The posterior-separable static information cost function derived from the neighborhood-based generalized entropy can then be written as

$$\begin{aligned} C_N^{static}(p,q;S) &= \sum_{i\in\mathscr{I}(q)} c_i \bar{q}_i \sum_{s\in S} \bar{p}_{i,s} D_{KL}(q_{i,s}||q_i) \\ &= \sum_{s\in S} (e_s^T pq) D_N(q_s||q), \end{aligned}$$

or alternatively as

$$C_N^{static}(p,q;S) = \sum_{i\in\mathscr{I}(q)} c_i \sum_{x\in X_i} (e_x^T q) D_{KL}(pe_x||pE_i^T q_i).$$

Proof. See the appendix, section A.9. □

The fact that $\bar{k}_N(q)$ is the Hessian of a convex function $H_N(q)$ means that we can apply the sufficient condition (18) in order to verify that Condition 6 is satisfied by any cost function in this family. For a flow cost function of the form (20), the marginal cost of increasing the probability of a jump to an arbitrary posterior q' is given by the divergence

$$D_N^*(q'||q) = \sum_{i\in\mathscr{I}(q)} \bar{q}'_i D_i(q'_i||q_i).$$

Under our assumption that D_i is bounded below by $c_i D_{KL}$ for each i, it follows that

$$D_N^*(q'||q) \;\geq \sum_{i \in \mathscr{I}(q)} c_i \bar{q}_i' D_{KL}(q_i'||q_i) \;=\; D_N(q'||q),$$

and condition (18) is verified. Thus under the assumptions stated above, any neighborhood-based cost function satisfies all of Conditions 1-6 for a flow information-cost function. If for each i, D_i is a positive multiple of an α-divergence with $\alpha_i > 1$, then (18) holds with a strict inequality for any $q' \neq q$. In this case, the cost function satisfies the strong form of Condition 6, so that there is a strict preference for gradual learning.

Lemma 5 allows us to write the static rational inattention problem (Theorem 1) directly in terms of an optimization over choice probabilities $\{\pi_x\}$ so as to maximize

$$\sum_{x \in X} e_x^T q_0 \sum_{a \in A} e_a^T \pi_x u_{x,a} - \theta C(\{\pi_x\}_{x \in X}, q_0; A). \tag{22}$$

As discussed previously, in the special case in which there is only a single neighborhood, this is the standard rational inattention problem. The relevance of alternative assumptions about the neighborhood structure is illustrated by the following result.

Lemma 6. *Consider a rational inattention problem (22) with a neighborhood-based information-cost function, and let x, x' be two states with the property that (i) $u_{a,x} = u_{a,x'}$ for all actions $a \in A$, and (ii) 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, $\pi_x^* = \pi_{x'}^*$.*

Proof. The result follows directly from the problem in (22) and the alternative expression for the cost function in Lemma 5. □

5.1 An Application: Psychometric Functions

The significance of Lemma 6 can be seen if we consider the predictions of rational inattention for a standard form of perceptual discrimination experiment. Suppose that the different states $X = \{1, 2, \ldots, N\}$ 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 $x > (N+1)/2$. 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. In such experiments, 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 $x > (N+1)/2$ or if $a = L$ and $x < (N+1)/2$, while $u_{x,a} = 0$ in all other cases. We shall assume that each of the possible stimuli is presented with equal prior probability, and hence (assuming that N is odd) that both responses have an equal ex ante probability of being correct.

The standard theory of rational inattention, in which the static information cost is mutual information, corresponds to a special case of a neighborhood-based cost function, in which all states belong to the unique neighborhood. Hence condition (ii) of Lemma 6 holds for any pair of states. Lemma 6 thus implies that if any two states result in the same payoff regardless of the action chosen, the frequency with which different actions will be chosen under an optimal policy must be the same in the two states.

In the problem just posed, this implies that the probability of response R must be the same for all states $x < (N+1)/2$, and also the same (but higher) for all states $x > (N+1)/2$. Changing the severity of the information constraint changes the degree to which the probability of responding R is higher when $x > (N+1)/2$, but it cannot change the prediction that the response probabilities should depend only on whether x is greater or less than $(N+1)/2$. This is illustrated in 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 $N = 20$.

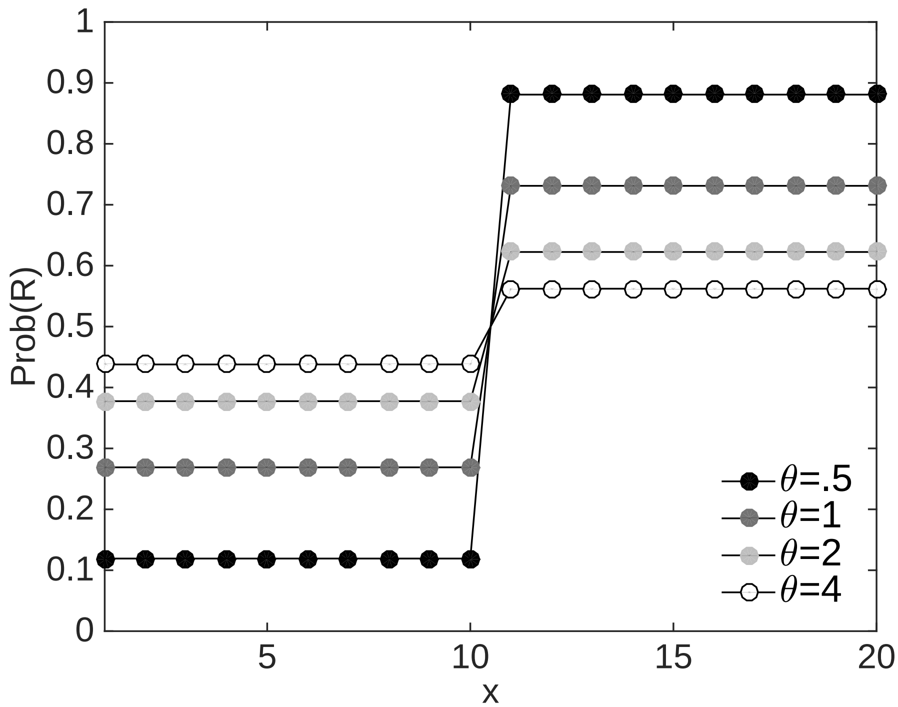

Figure 1: Predicted response probabilities with a mutual-information cost function, for alternative values of the cost parameter θ.

Alternatively, consider a posterior-separable neighborhood-based cost function in which the neighborhoods are given by

$$X_i = \{x_i, x_{i+1}\} \tag{23}$$

for $i = 1, 2, \ldots, N-1$. Thus two states belong to a common neighborhood if and only if they are either identical or 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.[33] These assumptions suffice to completely determine a static information cost function (Lemma 5).

With this alternative neighborhood structure, Lemma 6 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 Figure 2, which again shows the optimal response frequencies as a function of x, for alternative values of θ, in the case of the alternative neighborhood structure (23). 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.[34]

[33] If 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.

[34] For 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 a field of moving dots as to which of two opposing directions is the dominant direction of motion), 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 in this experiment 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.

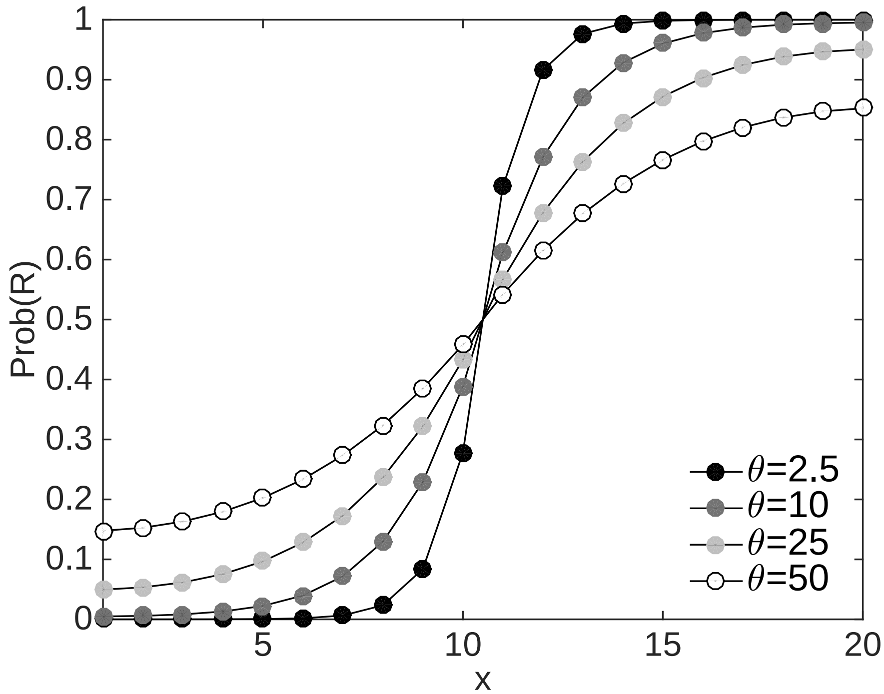

Figure 2: Predicted response probabilities with a neighborhood-based cost function, in which each neighborhood consists only of two adjacent states.

The continuity of choice probabilities across points at which there are discrete changes in payoffs is also an important issue for the global games literature (Morris and Yang (2016)). However, this literature typically assumes a continuum of states, and many of the perceptual experiments that we have just referred to are naturally modeled with a continuum of states as well. In the next sub-section, we consider a continuous-state limit of the example just analyzed, that can be used as a model of imprecise perception in such examples.

5.2 The Fisher-Information Cost Function

In this subsection, we continue our discussion of the neighborhood-based cost function proposed in the previous subsection, and consider the limit as the number of states of the world, $|X|$, becomes infinite. This example is motivated by the work of Yang (2015) and Morris and Yang (2016), who study global games (e.g. Morris and Shin (2001)) with endogenous information acquisition. However, we derive our limiting result for arbitrary action spaces and utility functions.

The result corresponds to a static rational inattention problem with a continuum of states, in which the information cost function is given by the average value of the Fisher information, a measure of the precision with which an information structure allows nearby states to be distinguished from each other (Cover and Thomas (2012)). Like Sims' proposal of a cost function proportional to Shannon's 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. But unlike Shannon's mutual information, our measure of the informativeness of an information structure based on the Fisher information depends on the topological structure of the state space.

This is of considerable significance for the literature on global games. In the well-known analysis of Morris and Shin (2001), 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). Instead 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. However, our results in the previous subsection cannot be applied directly to the model of Morris and Yang (2016), because the global games model in that paper assumes a continuum of states, whereas our analysis above supposes that $|X|$ is finite. To bridge this gap, we study an example of the static model implied by Theorem 1 with a particular neighborhood-based cost function, and consider the limit as the number of states becomes unboundedly large. We show that the example model converges to a static rational inattention model with a particular cost function, similar in certain respects to the leading example of Morris and Yang (2016), that satisfies the continuous choice condition established by those authors.

For each of a sequence of values for the finite integer N, we assume a neighborhood structure of the kind discussed in the previous subsection for a model with $N+1$ states. The set of states is ordered, $X^N = \{0, 1, \ldots, N\}$, and each pair of adjacent states forms a neighborhood, $X_j = \{i, i+1\}$, for all $j \in \{0, 1, \ldots, N-1\}$. We will also assume that there is an $N+1$st neighborhood containing all of the states. Note that N indexes both the number of states and the number of neighborhoods, which is always equal to the number of states. We consider the limit as $N \to \infty$.

To study this limit, we need to define how the initial beliefs, q_N, and the magnitude of the information costs vary with N. For the initial beliefs, we shall assume that there is a differentiable probability density function $f : [0,1] \to \mathbb{R}^+$, with full support on $[0,1]$, with a derivative that is Lipschitz continuous. Using this function, we define, for any $i \in X^N$,

$$e_i^T q_N = \int_{\frac{i}{N+1}}^{\frac{i+1}{N+1}} f(x)dx.$$

That is, for each value of N, the prior q_N is assumed to be a discrete approximation to the Lipschitz-continuous p.d.f. $f(x)$, which becomes increasingly accurate as $N \to \infty$.

For our neighborhood structures, we assume that that the constants associated with the cost of each neighborhood, c_j, are equal to N^2 for all $j < N$, and N^{-1} for $j = N$. In this particular example, the scaling ensures that the DM is neither able to determine the state with certainty, nor prevented from gathering any useful information, even as N is made arbitrarily large; moreover, the scaling ensures that the neighborhood containing all states plays no role in the limiting behavior, so that in the limit all information costs are local. We also scale the entire cost function by a constant, $\bar{\theta} > 0$.

We also need to define the set of actions, and the utility from those actions. We will assume the set of actions, A, remains fixed as N grows, and define the utility from a particular action, in a particular state, as

$$e_i^T u_{a,N} = \frac{\int_{\frac{i}{N+1}}^{\frac{i+1}{N+1}} f(x) u_a(x) dx}{e_i^T q_N}.$$

Here, the utility $u_a : [0,1] \to \mathbb{R}$ is a bounded measurable function for each action $a \in A$.[35] In other words, as N grows large, the prior converges to $f(x)$ and the utilities converge to the functions $u_a(x)$.

Under these assumptions, the static model of Theorem 1 can be written as

$$V_N(q_N;\bar{\theta}) = \max_{\pi_N \in \mathscr{P}(A), \{q_{a,N} \in \mathscr{P}(X^N)\}_{a\in A}} \sum_{a\in A} \pi_N(a)(u_{a,N}^T \cdot q_{a,N}) - \bar{\theta} \sum_{a\in A} \pi_N(a) D_N(q_{a,N}||q_N), \tag{24}$$

[35]Note that we do not require the payoff resulting from a action to be a continuous function of x at all points, though it will be continuous almost everywhere. This allows for the possibility that a DM's payoffs change discontinuously when the state x crosses some threshold, as in the kind of equilibria discussed by Yang (2015).

subject to the constraint that

$$\sum_{a\in A}\pi_N(a)q_{a,N}=q_N.$$

Here D_N denotes the divergence associated with the neighborhood-based cost function introduced above, specialized to the particular neighborhood structure of this section:

$$D_N(q_{a,N}||q_N)=N^2\sum_{j\in\mathbf{X}^N\setminus\{N\}}\bar{q}_{j,a,N}D_{KL}(q_{j,a,N}||q_{j,N})+N^{-1}D_{KL}(q_{a,N}||q_N).$$

The following theorem shows that the solution to this problem, both in terms of the value function and the optimal policies, converges to the solution of a static rational inattention problem with a continuous state space.

Theorem 3. *Consider the sequence of finite-state-space static rational inattention problems (24), with progressively larger state spaces indexed by the natural numbers N. Then there exists a sub-sequence of integers $n\in\mathbb{N}$ for which the solutions to the sub-sequence of problems converge, in the sense that*

i) $\lim_{n\to\infty}V_n(q_n;\bar{\theta})=V(q;\bar{\theta})$;

ii) $\lim_{n\to\infty}\pi_n^*=\pi^*$; *and*

iii) for all $a\in A$ and all $x\in[0,1]$, $\lim_{n\to\infty}\sum_{i=0}^{\lfloor xn\rfloor}e_i^T q_{a,n}^*=\int_0^x f_a^*(y)dy$.

Moreover, the limiting value function $V(q;\bar{\theta})$ is the value function for the following continuous-state-space static rational inattention problem:

$$V(f;\bar{\theta})=\sup_{\pi\in\mathscr{P}(A),\{f_a\in\mathscr{P}_{LipG}([0,1])\}_{a\in A}}\sum_{a\in A}\pi(a)\int_0^1 u_a(x)f_a(x)dx$$
$$-\frac{\bar{\theta}}{4}\sum_{a\in A}\{\pi(a)\int_0^1\frac{(f_a'(x))^2}{f_a(x)}dx\}+\frac{\bar{\theta}}{4}\int_0^1\frac{(f'(x))^2}{f(x)}dx,$$

subject to the constraint that, for all $x \in [0,1]$,

$$\sum_{a \in A} \pi(a) f_a(x) = f(x), \tag{25}$$

and where $\mathscr{P}_{LipG}([0,1])$ *denotes the set of differentiable probability density functions with full support on* $[0,1]$*, whose derivatives are Lipschitz-continuous. Furthermore, the limiting action probabilities* $\pi^*(a)$ *and posteriors* f_a^* *are the optimal policies for the continuous-state-space problem.*

Proof. See the appendix, section A.11. □

The static rational inattention problem for the limiting case of a continuous state space can be given an alternative, equivalent formulation, in which the objects of choice are the conditional probabilities of taking different actions in the different possible states, rather than the posteriors associated with different actions.

Lemma 7. *Consider the alternative continuous-state-space static rational inattention problem:*

$$\bar{V}(f;\bar{\theta}) \;=\; \sup_{p \in \mathscr{P}_{LipG}(A)} \int_0^1 f(x) \sum_{a \in A} p_a(x) u_a(x) dx - \frac{\bar{\theta}}{4} \int_0^1 f(x) I^{Fisher}(x;p) dx,$$

where $\mathscr{P}_{LipG}(A)$ *is the set of mappings* $p : [0,1] \to \mathscr{P}(A)$ *such that for each action a, the function* $p_a(x)$[36] *is a differentiable function of x with a Lipschitz-continuous derivative, and for any information structure* $p \in \mathscr{P}_{LipG}(A)$*, the* Fisher information *at state* $x \in X$ *is defined as*

$$I^{Fisher}(x;p) \;\equiv\; \sum_{a \in A} \frac{(p'_a(x))^2}{p_a(x)}.$$

[36]Here for any $x \in [0,1]$, we use the notation $p_a(x)$ to indicate the probability of action a implied by the probability distribution $p(x) \in \mathscr{P}(A)$.

This problem is equivalent to the one defined in Theorem 3, in the sense that the information structure p^* that solves this problem defines action probabilities and posteriors

$$\pi^*(a) = \int_0^1 f(x) p_a^*(x), \qquad f_a^*(x) = \frac{f(x) p_a^*(x)}{\pi^*(a)} \tag{26}$$

that solve the problem in Theorem 3, and conversely, the action probabilities and posteriors $\{\pi^*(a), f_a^*\}$ that solve the problem stated in the theorem define state-contingent action probabilities

$$p_a^*(x) = \frac{\pi^*(a) f_a^*(x)}{f(x)} \tag{27}$$

that solve the problem stated here. Moreover, the maximum achievable value is the same for both problems: $\bar{V}(f;\bar{\theta}) = V(f;\bar{\theta})$.

Proof. See the appendix, section A.12. □

Theorem 3 shows that the limit of our neighborhoods problem converges to a static rational inattention problem with a particular cost function. That cost function is just the expected value of the Fisher information $I^{Fisher}(x;p)$, defined locally for each element of the continuum of possible states x, with the expectation taken with respect to the prior over possible states.[37] This cost function, unlike mutual information, depends only on the degree to which the information structure allows states to be distinguished from ones extremely close to them (under the topology of the real line); and unlike the rational inattention problem based on mutual information, this static mutual information problem will generate the smoothness of responses across discrete changes in payoffs shown in Figure 2.

For these reasons, we believe that the Fisher-information cost function is likely to be more appropriate than mutual information in a wide range of settings. It should also be noted that, as in the case of Sims' theory of rational inattention, the Fisher-information cost

[37] This aggregate Fisher information has also proven useful in a variety of physics applications (Frieden (2004)).

function has only a single degree of freedom. We thus obtain a rational inattention theory for problems with a continuous space that yields highly specific predictions, albeit different ones from Sims' theory.

We can apply this result to the problem considered in Morris and Yang (2016). 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 $\bar{\theta} > 0$ and all parameterizations of the relevant utility function, $s(x)$ is absolutely continuous. Our Theorem 3 provides an example of more primitive assumptions that would guarantee continuous choice in this sense.

We believe that these results show the usefulness of our continuous-time model of evidence accumulation as a micro-foundation for interesting classes of static rational-inattention problems, with properties that are relevant for economic applications. It remains for us to explain the justification for our proposed formulation of the continuous-time model of evidence accumulation itself.

6 Derivation of the Continuous-Time Model

We now show how the continuous-time model proposed in section 2 can be obtained as the limit of a discrete-time model of sequential evidence accumulation, with a sequence of endogenous signals as in dynamic rational inattention models like that of Steiner et al. (2017), and an information cost function for each of the individual signals that satisfies the properties proposed for flow cost functions in section 4. In particular, we will justify the link proposed above between a second-order approximation to the information cost function for an individual signal and the information-cost matrix function defined in section 2.

We study a dynamic problem in which the DM has repeated opportunities to gather information before making a decision. The state of the world, $x \in X$, remains constant over time. At each time t, the DM can either stop and take an action $a \in A$, or continue and receive a signal drawn from the information structure $\{p_{t,x} \in \mathscr{P}(S)\}_{x \in X}$, for some signal alphabet S. We assume that the number of potential actions is weakly less than the number of states, $|A| \leq |X|$.

We also assume that the signal alphabet S is finite and fixed over time, with $|S| \geq 2|X|+1$. However, the information structure $\{p_{t,x}\}_{x \in X}$ is a choice variable that can be state- and time-dependent. Fixing the signal alphabet S has no economic meaning, because the information content of receiving a particular signal $s \in S$ can change between periods. The assumption allows us to assume a finite information structure and invoke the results from section 4.[38] As a technical device, we assume that S contains one signal, $\bar{s}$, that is required to be uninformative. This assumption is a technical device to ensure that the DM can choose to mix any arbitrary signal structure with an uninformative one, even if she has already used up her "useful" signals.

The DM's prior beliefs at time t, before receiving the signal, are denoted q_t. Each time period has a length Δ. Let τ denote the time at which the DM stops and makes a decision, with $\tau = 0$ corresponding to making a decision without acquiring any information. At this time, the DM receives utility $u(x,a) - \kappa\tau$ if she takes action a at time τ and the true state of the world is x. As in the previous sections, let $\hat{u}(q_\tau)$ be the utility (not including the penalty for delay) associated with taking an optimal action under beliefs q_τ. The parameter κ governs the size of the penalty the DM faces from delaying his decision. The reason the DM does not make a decision immediately is that she is able to gather information, and make a more-informed decision. The setup thus far is essentially identical to the continuous-time

[38]As mentioned previously, the work of Ay et al. (2014) discusses how to extend the Chentsov (1982) theorems to infinite-dimensional structures. We conjecture that their results would allow us to extend our theorems to infinite signal spaces, but do not attempt such an extension here.

model described previously.

The DM can choose an information structure that depends on the current time and past history of the signals received. As we will see, the problem has a Markov structure, and the current time's "prior," q_t, summarizes all of the relevant information that the DM needs to design the information structure. The DM is constrained to satisfy

$$E_0[\frac{\Delta}{\rho}\sum_{j=0}^{\tau\Delta^{-1}-1} C(\{p_{\Delta j}, q_{\Delta j}; S)^{\rho}]^{\frac{1}{\rho}} \leq \Delta c E_0[\tau], \tag{28}$$

if the DM choose to acquire any information at all ($\tau > 0$ always in this case). In words, the L^{ρ}-norm of the flow information cost function $C(\cdot)$ over time and possible histories must be less that the constant c per unit time.

In the limit as $\rho \to \infty$, this would approach a per-period constraint on the amount of information the DM can obtain. For finite values of ρ, the DM can allocate more information gathering to states and times in which it is more advantageous to gather more information. We assume, however, that $\rho > 1$, for reasons that we discuss later. We also assume that the flow cost function $C(\cdot)$ satisfies Conditions 1-6 stated in section 4.

Let $V(q_0;\Delta)$ denote the value of the solution to the sequence problem for a DM with prior beliefs q_0, and let q_τ denote the DM's beliefs when stopping to make a decision. The DM's problem is

$$V(q_0;\Delta) = \max_{\{p_{\Delta j}\},\tau} E_0[\hat{u}(q_\tau) - \kappa\tau)],$$

subject to the information-cost constraint (28). The dual version of this problem can be written as

$$W(q_0,\lambda;\Delta) = \max_{\{p_{\Delta j}\},\tau} E_0[\hat{u}(q_\tau) - \kappa\tau)] - \\ \lambda E_0[\Delta^{1-\rho} \sum_{j=0}^{\tau\Delta^{-1}-1} \{\frac{1}{\rho} C(p_{\Delta j}, q_{\Delta j}; S)^\rho - \Delta^\rho c^\rho\}]. \tag{29}$$

Here, the function $W(q_0,\lambda;\Delta)$ can be thought of as the value function of a different problem, in which there is a cost of gathering information proportional to $\lambda\frac{1}{\rho}C(\cdot)^\rho$. In what follows, we will refer to the function W as the value function, bearing in mind that λ is not actually exogenous to the problem. We will proceed under the assumption that $\lambda \in (0, \kappa c^{-\rho})$. In our proofs, we demonstrate that there is no duality gap in the continuous time limit of this problem, and that our assumption about λ is without loss of generality.

We begin by describing the recursive representation for the value function $W(q_t,\lambda;\Delta)$, and discussing certain technical lemmas that are necessary to establish our main results. The value function has a recursive representation:

$$\begin{aligned} W(q_t,\lambda;\Delta) &= \max\{\max_{p_t} -\kappa\Delta + \lambda\Delta^{1-\rho}(\Delta^\rho c^\rho - \frac{1}{\rho}C(p_t,q_t;S)^\rho) + \\ &\quad \sum_{s\in S}(e_s^T p_t q_t) W(q_{t+\Delta,s},\lambda;\Delta),\ \hat{u}(q_t)\}, \end{aligned}$$

where $q_{t+\Delta,s}$ is pinned down by Bayes' rule. In standard rational inattention problems, it is without loss of generality to equate signals and actions. In this problem, when the DM does not stop and make a decision, the "action" is updating one's beliefs. Rather than consider a probability distribution over signals, and then an updating of beliefs by Bayes' rule, one can consider the DM to be choosing a probability distribution over posteriors, subject to the constraint that the expectation of the posterior is equal to the prior.[39]

[39]The notion of choosing a probability distribution over posteriors appears in Kamenica and Gentzkow (2011), Caplin and Dean (2015), and Caplin et al. (2017), among other papers.

To begin our analysis, we note that the value function $W(q_t,\lambda;\Delta)$ is well-behaved:

Lemma 8. *The value function $W(q_t,\lambda;\Delta)$ is bounded on $q_t \in \mathscr{P}(X)$, and convex in q. The optimal stopping time τ_Δ is bounded in expectation by a constant, $\bar{\tau}$, for all Δ:*

$$E_0[\tau_\Delta] \leq \bar{\tau}.$$

Proof. See the appendix, section A.13. □

The boundedness of the value function follows from the setup of the problem: ultimately, the DM will make a decision, and the utility from making the best possible decision in the best possible state of the world is finite. The convexity of the value function is what motivates the DM to acquire information. By updating her beliefs from q to either q' or q'', with $q = \alpha q'' + (1-\alpha)q'$ for some $\alpha \in (0,1)$, the DM improves her welfare by enabling better decision making. That the optimal stopping time is bounded in expectation follows from an obvious point: waiting too long to make a decision will eventually become worse, even if the DM eventually makes the best possible decision, than making the worst possible decision immediately.

Next, we show that, because of the curvature (ρ) that we impose, the DM will choose, under any optimal policy, to gather only a small amount of information in each time period, as the length of each time period shrinks.

Lemma 9. *Let $n \in \mathbb{N}$ denote a sequence such that $\lim_{n\to\infty}\Delta_n = 0$. Any associated sequence of optimal policies $p^*_{t,n}$ satisfies, for all elements of the sequence,*

$$C(p^*_{t,n},q_{t,n};S) \leq \left(\frac{\theta}{\lambda}\right)^{\frac{1}{\rho-1}}\Delta_n,$$

where $\theta = \lambda\left(\rho\frac{\kappa-\lambda c^\rho}{\lambda(\rho-1)}\right)^{\frac{\rho-1}{\rho}}$.

Proof. See appendix, section A.14. □

The key step in proving this lemma is demonstrating that, as the time period shrinks, the optimal quantity of information acquired vanishes at a sufficiently fast rate. The convergence of the information structure to an uninformative one, as the time period shrinks, allows us to use the approximation described in Theorem 2 to study the continuous-time limit of the sequential evidence accumulation model. The assumption that $\rho > 1$ is critical to generating this result. When $\rho = 1$, the DM has no particular desire to smooth the quantity of information gathered over time, and might choose to gather a large quantity of information in a single period (as in Steiner et al. (2017)).

We next discuss the convergence of an arbitrary sequence of stochastic processes for beliefs (denoted $q_{t,m}$) and of stopping times (denoted τ_m) to their continuous-time limits, under the assumption that the policies generating them satisfy the bound in Lemma 9 and the bound on expected stopping times. This lemma applies to a sequence of optimal policies, but also to sequences of sub-optimal policies. The lemma describes the convergence of the beliefs process to a martingale, which is not necessarily a diffusion (it may have jumps, or even be a semi-martingale that is not a jump-diffusion).

Lemma 10. *Let Δ_m, $m \in \mathbb{N}$, denote a sequence such that $\lim_{m\to\infty} \Delta_m = 0$. Let $p_m(q)$ denote a sequence of Markov policies satisfying the bound in Lemma 9. Let $q_{t,m}$ denote the stochastic process for the DM's beliefs at time t, under such a policy, and let τ_m be a sequence of stopping policies such that $E_0[\tau_m] \leq \bar{\tau}$.*

There exists a sub-sequence $n \in \mathbb{N}$ and a probability space such that:

i) The beliefs $q_{t,n}$ and the stopping time τ_n converge almost surely to a martingale q_t and a stopping time τ.

ii) The martingale q_t can be represented in terms of its semi-martingale characteristics,

$$B_t = -\int_0^t (\int_{\mathbb{R}^{|X|}\setminus\{0\}} \psi_s(x) x dx) dA_s$$

$$C_t = \int_0^t D(q_{s^-})\sigma_s \sigma_s^T D(q_{s^-}) dA_s$$

$$\nu_t(x) = dA_t \psi_t(x),$$

where σ_s is an $|X| \times |X|$ matrix-valued predictable stochastic process, satisfying $q_{s^-}^T \sigma_s = \vec{0}$, ψ_s is a measure on $\mathbb{R}^{|X|} \setminus \{0\}$ such that $q_{s^-} + x \in \mathscr{P}(X)$ and $q_{s^-} + x \ll q_{s^-}$ for all x in the support of ψ_s, and dA_s is the increment of a weakly increasing process.

iii) For all stopping times T,

$$E_t[\int_t^T \{\frac{1}{2} tr[\sigma_s \sigma_s^T k(q_{s^-})] + \int_{\mathbb{R}^{|X|}\setminus\{0\}} \psi_s(x) D^*(q_{s^-} + x || q_{s^-}) dx\} dA_s] \leq (\frac{\theta}{\lambda})^{\frac{1}{\rho-1}} E_t[T - t].$$

iv) The limit of the cumulative information cost is bounded below,

$$\lim_{n\to\infty} E_0[\int_0^{\tau_n} \Delta_n^{1-\rho} C(p_n(q_{t,n}), q_{t,n}; S)^\rho dt] \geq E_t[\int_0^\tau \{\frac{1}{2} tr[\sigma_s \sigma_s^T k(q_{s^-})] + \int_{\mathbb{R}^{|X|}\setminus\{0\}} \psi_s(x) D^*(q_{s^-} + x || q_{s^-}) dx\}^\rho (\frac{dA_s}{ds})^\rho ds].$$

Proof. See the appendix, section A.15. □

In essence, the stochastic process $q_{t,n}$ converges to a jump-diffusion process. The semi-martingale characteristics, B_t, C_t, ν_t, summarize the DM's policy function. They have a representation as a function of σ_t, ψ_t, A_t because of the need for beliefs to remain the simplex, and the property that, once a state $x \in X$ has been assigned zero probability, it will be assigned zero probability forever after.

To finish the proof, we resolve several issues. We show that the constraint given in Lemma 9 binds. We show that the limiting value function W is unique, that duality holds ($V = W$ for a suitable choice of λ), and that the the limit of V is the solution to the continuous-time problem described in section 2. We also show that there is a sequence of (possibly sub-optimal) policies in the discrete-time model that achieve, in the limit, the optimal utility and converge to a diffusion. Moreover, if the cost function $C(p,q;S)$ exhibits a strict preference for gradual learning (it satisfies Condition 6 strictly for $q' \neq q$), then all sequences of optimal policies converge to diffusions that are optimal policies of the continuous-time model.

Theorem 4. *Let $n \in \mathbb{N}$ index a sub-sequence of policies described in Lemma 10. There exists a $\lambda^* \in (0, \kappa c^{-\rho})$ such that*

$$\lim_{n\to\infty} W(q_t, \lambda^*; \Delta_n) = \lim_{n\to\infty} V(q_0; \Delta_n) = V(q_0),$$

where $V(q_0)$ is the solution to the continuous-time problem described in section 2, with $\chi = \rho^{\rho^{-1}} c$ and $\mu = \kappa$. There exists a sequence of policies in the discrete-time models that achieve, in the limit, the value function $V(q_0)$ and for which the associated belief process, $q_{t,n}$, and stopping time τ_n converges in law to a belief process q_t^* and stopping time τ^* that are induced by an optimal policy in the continuous-time model (and hence q_t^* is a diffusion). If the cost function exhibits a strict preference for gradual learning, every convergent sub-sequence of belief processes $q_{t,n}^*$ associated with optimal policies in the discrete-time model converges in law to a diffusion.

Proof. See the appendix, section A.16. □

We have shown that the DM's behavior in the continuous-time problem can be thought of as an approximation of her behavior in discrete-time problems with flow cost functions drawn from a very general class. These convergence results can be viewed as offering a sort of micro-foundation for the continuous-time model, and in particular for our assumptions in section 2 about the information-cost matrix function.

7 Conclusion

We have derived a continuous-time rational-inattention model as the limit of a discrete-time sequential evidence accumulation problem. While assumptions about the cost of more precise signals in the sequential evidence accumulation problem are important determinants of the predictions obtained for the choices that will be made, we have shown that relatively specific conclusions are possible about the endogenous information structure in the continuous-time limit, even under relatively general assumptions about the flow cost function for individual signals. This is because in the limit of a very large number of successive signals, each of which is only minimally informative, only the local properties of the flow cost function near purely uninformative information structures matter. The relevant properties of the flow cost function can be summarized by a matrix-valued function defined on the space of possible posterior beliefs, that we call the information-cost matrix function. This summarizes the degree to which it is costly to further distinguish between different pairs of possible states, when one's posterior belief given observations to that point is a particular point in the probability simplex. We view it as desirable that our framework retains the flexibility to allow different specifications of the degree to which it is intrinsically difficult to distinguish certain pairs of states from one another.

Quite generally, the solution to our continuous-time rational inattention model can be characterized by the solution to a partial differential equation (HJB equation) that involves the information-cost matrix function. For a broad class of possible specifications of the information-cost function, we are able to solve this equation, and further show that the solution to our continuous-time dynamic model is equivalent to the solution to a static rational-inattention problem, with a particular posterior-separable information-cost function (which depends on the information-cost matrix function, and hence on the local properties of the flow information-cost function). The use of posterior-separable cost functions of this kind in static rational-inattention problems can thus be justified summarizing the implications of a dynamic evidence accumulation process.

Among the static cost functions that can be justified in this way is the mutual-information cost function proposed by Sims, but we show that it is not the only static cost function that can be given such a justification. We give particular attention to the existence of cost functions that can be justified as the outcome of a dynamic evidence accumulation process, but that incorporate an assumption that "nearby" states are more difficult to distinguish from one another, unlike the mutual-information cost function. We exhibit a class of "neighborhood-based" cost functions that are convenient specifications of this kind, and discuss in particular a limiting case of such functions that can be used in problems with a continuous state space, our Fisher-information cost function. This cost function has the attractive feature that an optimal endogenous information structure implies that action probabilities will vary continuously with the state, even when action payoffs jump discontinuously. A model of this kind better matches observed behavior in perceptual experiments, and we suspect that it represents a more realistic assumption for economic applications as well, such as global game models of the kind analyzed by Yang (2015) and Morris and Yang (2016).

References

Shun-ichi Amari and Hiroshi Nagaoka. *Methods of information geometry*, volume 191. American Mathematical Soc., 2007.

Nihat Ay, Jürgen Jost, Hông Vân Lê, and Lorenz Schwachhöfer. Information geometry and sufficient statistics. *Probability Theory and Related Fields*, pages 1–38, 2014.

David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, 24(2):265–272, 1953.

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

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Unpublished manuscript*, July 2017.

Yeon-Koo Che and Konrad Mierendorff. Optimal sequential decision with limited attention. *Unpublished manuscript*, 2016.

Nikolai Nikolaevich Chentsov. *Statistical decision rules and optimal inference*. Number 53. American Mathematical Soc., 1982.

John A. Clithero. Improving out-of-sample predictions using response times and a model of the decision process. *Unpublished manuscript*, June 2016.

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.

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.

Mogens Fosgerau, Emerson Melo, and Matthew Shum. Discrete choice and rational inattention: A general equivalence result. *Unpublished manuscript*, 2016.

B Roy Frieden. *Science from Fisher information: a unification*. Cambridge University Press, 2004.

Drew Fudenberg, Philipp Strack, and Tomasz Strzalecki. Stochastic choice and optimal sequential sampling. *Unpublished manuscript*, 2015.

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. *Unpublished manuscript*, 2014.

Frank Huettner, Tamer Boyaci, and Yalcin Akcay. Consumer choice under limited attention when options have different information costs. *Unpublished manuscript*, 2016.

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.

Filip Matêjka, 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. Global games: theory and applications. 2001.

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.

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.

Christopher A Sims. Rational inattention and monetary economics. *Handbook of Monetary Economics*, 3:155–181, 2010.

Jakub Steiner, Colin Stewart, and Filip Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.

Satohiro Tajima, Jan Drugowitsch, and Alexandre Pouget. Optimal policy for value-based decision-making. *Nature communications*, 7, 2016.

Tim van Erven and Peter Harremoës. Rényi divergence and Kullback-Leibler divergence. *IEEE Transactions on Information Theory*, 60:3797–3820, 2014.

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.

Weijie Zhong. Optimal dynamic information acquisition. *Unpublished manuscript*, March 2017.