Algorithmic Choice Architecture for Boundedly Rational Consumers

Stefan Bucher and Peter Dayan^{*}

This version: November 2, 2023

Updated regularly, please click here for the latest version.

Abstract

Choice architecture and recommender systems both address information overload but have developed largely independently of each other and make strong assumptions about decisionmakers' unobserved preferences. In this paper, we introduce cognitive information filters as an algorithmic approach to choice architecture that mitigates information overload in a more principled and effective manner: our method combines machine learning with a cognitive model of choice behavior to solve the economic problem of nudging or persuading decision-makers by tailoring information to their revealed preferences and cognitive constraints. We first develop a rational-inattention model of multi-attribute choice to describe the behavior of a consumer (receiver) facing information costs. We then use reinforcement learning to solve the information design problem of a sender choosing which options and attributes are accessible to the receiver. Observing only the receiver's choices, the sender learns from repeated interactions which information is most effective in attaining desirable behavioral outcomes. By inferring preferences from boundedly rational behavior, our methodology can optimize for revealed welfare and hence promises to be (1) less paternalistic than traditional nudging and (2) less susceptible to misalignment than recommender systems optimizing for imperfect welfare proxies such as engagement. This has implications beyond economics and marketing, for example for digital platforms and alignment research in artificial intelligence.

JEL classification: C90, D11, D82, D83, D91, M31

Keywords: choice architecture, information design, rational inattention, reinforcement learn-

ing

^{*}Tübingen AI Center & Max Planck Institute for Biological Cybernetics. Contact: stefan.bucher@nyu.edu. The authors thank Adam Brandenburger, Andrew Caplin, Sam Kapon, and Oleg Solopchuk for helpful discussions, and anonymous reviewers for the MIT Conference on Digital Experimentation and the NeurIPS Workshop on Information-Theoretic Principles in Cognitive Systems for helpful comments.

"It's Not Information Overload. It's Filter Failure."

- Clay Shirky

1 Introduction

Information overload is ubiquitous because constraints on our cognitive capacity to process information adversely impact the quality of the decisions we make.¹ Limited attention and cognition have important consequences for welfare and markets (McFadden, 2023), and have consequently become of central interest to economists and policymakers concerned that "individuals are able to pay only limited attention to important aspects of their environment, often have a difficult time processing information, and make cognitive errors even in simple situations," as stated in the National Academies' recent consensus study report on behavioral economics (Buttenheim et al., 2023, p. 7).

Since the field's inception, behavioral economists have thus devised remedies such as choice architecture and "nudging," with the goal of helping policymakers induce desirable behavioral outcomes (Thaler and Sunstein, 2008; Johnson, 2021). Because attention is a scarce good in the digital economy, unsurprisingly, firms are also competing for it with increasingly sophisticated recommender systems affecting our choices and beliefs (Fleder and Hosanagar, 2009; Aridor et al., 2022). Yet, identifying reliable nudges with significant effect sizes can be difficult and costly (DellaVigna and Linos, 2022; Mertens et al., 2022; Maier et al., 2022) – particularly with heterogeneous populations. By contrast, algorithmic recommender systems are designed to learn from observing individual user choices, but they often suffer from misalignment due to the difficulty of inferring users' preferences from their boundedly rational behavior (Kleinberg et al., 2022; Hébert and Zhong, 2022; McLaughlin and Spiess, 2022).

This paper addresses the question of whether and how a choice architect can learn to provide a boundedly rational decision-maker with the information that most effectively induces desirable behavioral outcomes, by tailoring it to their revealed preferences and cognitive constraints. We achieve a solution by using machine learning based on a model of stochastic choice under cognitive costs. By combining the advantages of recommender systems with careful revealed preference, our principled approach to algorithmic choice architecture

¹The notion that information overload is detrimental to decision-making is often attributed to management scholars (e.g. Jacoby et al., 1974; Jacoby, 1984), but the concept itself goes back much further (Blair, 2011).

promises to be more effective and less paternalistic, in allowing to explicitly maximize consumer welfare.

We introduce the concept of **cognitive information filters**, which we define as a method selecting a subset of information to be shown to a decision-maker to affect their choice behavior by changing their cognitive load. Specifically, we consider a setting in which a consumer (henceforth *receiver*) repeatedly makes multi-attribute choices. Confronted with information about a potentially large number of attributes (features) of a large number of options, the receiver is free to choose what information to attend to (at a cost) and what information to disregard. Observing only the receiver's choices, a choice architect (henceforth *sender*) can affect the receiver's choice behavior by determining which choice options and information on which attributes are accessible (in principle) to them, either to nudge them (with the goal of increasing their welfare) or persuade them. The sender's incentives may or may not be aligned with the receiver's.

The paper's contributions include a model of the receiver and a model of the sender, both of which are novel. To model the **receiver**, we introduce a model of multi-attribute search and choice under information costs. Our model builds on the rational-inattention literature (Sims, 2003; Mačkowiak et al., 2023), to which we contribute a multi-attribute choice model with an analytical solution. We do so by generalizing the result of Bucher and Caplin (2021) to identify conditions under which the (undistorted) multinomial logit choice probabilities of Matejka and McKay (2015) do not only satisfy their necessary conditions, but also the sufficient conditions of Caplin et al. (2019). The resulting discrete-choice model describes a consumer optimally acquiring costly information before making a multi-attribute choice, and promises to be a powerful alternative to models of costly sequential search in empirical (multi-attribute) settings in which the latter are intractable or empirically less relevant (cf. Honka and Chintagunta, 2017).

Our model allows for only a subset of information to be accessible, and makes a number of novel predictions regarding consumer demand and welfare under different information sets. In an example with quality and price as the only attributes, we show how the price elasticity of demand depends on the presence of information on quality: our model predicts demand is more elastic to the price in presence of information on quality than in its absence. This effect operates through an information channel based on the observation that positively correlated quality and price are informational complements: the consumer has a higher incentive to attend to the price in the presence of information on quality than in its absence.

The sender's problem is to choose an *information filter* determining which options and attributes are available to the receiver. The sender does not, however, know the receiver's type, which consists of their preferences and information costs. Instead, the sender infers the receiver's type in repeated interactions from observing their choices, using revealed preference. The sender thus solves a dynamic, model-based (inverse) reinforcement-learning problem, which we formulate as a partially observable Markov decision process (POMDP). POMDPs have originated in robotics (Kaelbling et al., 1998) and model sequential decisions under uncertainty in which *latent* states (in our case, the receiver's type) are inferred from noisy observations (in our case, the receiver's choices).

To solve the sender's information design problem, we rely on a model-based, online reinforcement learning and planning algorithm: Partially Observable Monte Carlo Planning (Silver and Veness, 2010) relies on Monte Carlo Tree Search (MCTS) to estimate continuation values, and on particle filters to update beliefs.² This algorithm makes an approximate solution feasible even if the state space is large. Using model-based reinforcement learning based on our structural receiver model has several important advantages: First, it is more robust to changes (i.e., less susceptible to the Lucas (1976) critique) than model-free reinforcement learning, and more explainable. Second, it is more data efficient because the exploration is more directed. And third, it allows inferring and optimizing for otherwise unobservable consumer welfare.

In simulations, we demonstrate the receiver indeed attains higher welfare under the filtered information than under the full-information benchmark.

These results illustrate our approach to algorithmic choice architecture promises to have several advantages over traditional nudging (Thaler and Sunstein, 2008). First, the approach is more principled in explicitly formulating the problem and constraints it is intended to address. Second, it is less paternalistic in that it infers consumer preferences rather than assuming them. Third, it is more personalized and can thus meet the needs of heterogeneous populations. These advantages also promise recommender systems that are better aligned with the preferences of boundedly rational users. Our findings have implications beyond economics and marketing, for example for alignment research in artificial intelligence.

The paper proceeds as follows: after discussing the related literature, section 2 introduces the model. Section 3 builds economic intuition by discussing a simple example in which the receiver type is observable. Section 4 provides the solution to the rational inattention model of multi-attribute choice characterizing the receiver's behavior. Section 5 formulates the sender problem, and demonstrates how to solve it using reinforcement learning algorithms. Section 6 concludes. Appendix A summarizes the notation; all proofs are given in Appendix B.

 $^{^2\}mathrm{MCTS}$ is also a crucial component of DeepMind's AlphaGo (Silver et al., 2017) and MuZero (Schrittwieser et al., 2020) algorithms.

Related Literature

This paper contributes to several strands of literature. First and foremost, our paper is a contribution to the emerging literature in cognitive economics, which seeks to understand the cognitive foundations of choice behavior (e.g. Woodford, 2020; Enke and Graeber, 2021; Frydman and Jin, 2022; Glimcher, 2022; Caplin, 2023). It also relates to models of costly sequential search in economics (Stigler, 1961; Weitzman, 1979; Santos et al., 2012) and marketing (Honka et al., 2019; Compiani et al., 2021; Ursu et al., 2023), as well as the literature on choice overload (Iyengar and Lepper, 2000; Scheibehenne et al., 2010; Chernev et al., 2015; Dean et al., 2022). Specifically, our paper focuses on information costs in multi-attribute choice, and thus complements recent accounts of multi-attribute choice with sequential information acquisition in economics (Sanjurjo, 2017) and psychology (Busemeyer et al., 2019; Callaway et al., 2022; Yang and Krajbich, 2023). Our choice to abstract from the dynamics of information acquisition is motivated by recent equivalence results expressing the choices resulting from a dynamic choice process in static terms (Choi et al., 2018; Hébert and Woodford, 2021). Our multi-attribute choice model is in the realm of multivariate rational inattention (Dewan, 2020; Miao et al., 2022), to which it contributes a simple analytical solution imposing only a weak exchangeability condition on prior beliefs.

Second, our sender solves an information design problem reminiscent of Bayesian persuasion (Kamenica and Gentzkow, 2011; Kamenica, 2019; Bergemann and Morris, 2019) with a rationally inattentive receiver (Bloedel and Segal, 2021). The main difference to Bloedel and Segal (2021) is that our sender does not know the receiver's preferences. Because our focus is on practical applications, we restrict the sender to disclosing truthful information to a strategically naïve receiver, instead of assuming that the sender has commitment power. Our paper also relates to a literature on algorithmic Bayesian persuasion (Dughmi and Xu, 2019) in economics and computation.

Third, our paper also contributes to the literature on website morphing (Hauser et al., 2009a,b, 2014), which models the multi-armed bandit problem of a website designer serving consumers information that matches their "cognitive type". In contrast to existing approaches, we explicitly model the receiver's cognitive constraints. The resulting POMDP is more general than the multi-armed bandits and hidden Markov models (Liberali and Ferecatu, 2022) of the website morphing literature. In the absence of a Gittins index solution, we solve the problem using modern reinforcement learning algorithms. Our paper thus also contributes to the recent use of reinforcement learning methods in marketing (Liu, 2023). In contrast to existing work (Liu, 2022), we use model-based reinforcement learning, whose reliance on the receiver model has the advantage of being more data-efficient, while also allowing us to explicitly optimize for consumer welfare.

2 Model

A sender (Alice) observes, in each period $t \in \{0, ..., \infty\}$, the choice of a receiver (Bob) from a grand set \mathcal{A} of m actions, which are characterized by a set \mathcal{F} of n attributes. The attribute values for all actions are encoded in an $m \times n$ matrix $\mathbf{X}_t \in \mathcal{X}$ whose rows correspond to actions and columns to attributes, and which is drawn, independently across periods, from a common prior

$$\mathbf{X}_t \stackrel{iid}{\sim} \mu \in \Delta(\mathcal{X})$$

The only assumption we impose on μ is that it is row-exchangeable.

Definition 1. $\mu \in \Delta(\mathcal{X})$ is row-exchangeable if, for any $m \times m$ permutation matrix \mathbf{P}_m ,

$$\mu(\mathbf{P}_m \mathbf{X}) = \mu(\mathbf{X}) \qquad \forall \mathbf{X} \in \mathcal{X}.$$
(1)

This assumption allows for arbitrary statistical dependency across attributes, and for values to be correlated across actions as long as they are exchangeable.³ For example, it could be the case that exactly one action has a given attribute, as long as each action is equally likely to be that one.

Alice's Choice In each period, Alice determines which actions and attributes are available to Bob by choosing an *information filter* consisting of an ordered action selection matrix $\mathbf{A}_t \in \mathcal{A} \subseteq \mathbb{R}^{k \times m}$ and an ordered feature selection matrix $\mathbf{F}_t \in \mathcal{F} \subseteq \mathbb{R}^{l \times n}$ (with $1 \le k \le m$ and $1 \le l \le n$ chosen by Alice). Both matrices must have entries in $\{0, 1\}$ with exactly one positive entry per row and at most one per column; with the interpretation that columns correspond to actions or attributes, respectively, and the rows determine which subset of – and order in which – these are presented.

Bob's Choice Bob is myopic and strategically naïve with respect to Alice, and only has access to the filtered state

$$\mathbf{Y}_t := \mathbf{A}_t \mathbf{X}_t \mathbf{F}_t^T \in \mathcal{Y}.$$

Bob's problem is to choose an option from choice set $A(\mathbf{A}) := \{a \in \mathcal{A} : \sum_i \mathbf{A}_{ia} > 0\}$, and his utility function is given by $u_{\theta} : \mathcal{A} \times \mathbb{R}^{k \times n} \to \mathbb{R}$ where $\theta = (\boldsymbol{w}, \kappa) \in \Theta$ is Bob's type, which consists of a preference parameter \boldsymbol{w} along with a marginal cost parameter κ to be used below. The only assumption we impose on Bob's utility function is that it is invariant under permutation, defined as follows.

³The assumption is satisfied, for instance, by any row-column-exchangeable (RCE; Aldous et al., 2006, p. 123) random matrix \mathbf{X} ; column-exchangeability however is not necessary.

Definition 2. The utility function $u_{\theta}(a, \mathbf{AX})$ is invariant under permutation if, for any \mathbf{A} and any $k \times k$ permutation matrix \mathbf{P}_k^{ϱ} with associated permutation $\varrho : A(\mathbf{A}) \to A(\mathbf{A})$ satisfying $\mathbf{e}_{\varrho(a)} = \mathbf{P}_k^{\varrho} \mathbf{e}_a$ for all $a \in A(\mathbf{A})$, it is the case that

$$u_{\theta}(\varrho(a), \mathbf{P}_{k}^{\varrho} \mathbf{A} \mathbf{X}) = u_{\theta}(a, \mathbf{A} \mathbf{X}) \qquad \forall a \in A(\mathbf{A}), \, \mathbf{X} \in \mathcal{X}.$$
⁽²⁾

Bob's Information Bob observes his type θ and Alice's choice of \mathbf{A}_t and \mathbf{F}_t , but not the state \mathbf{X}_t . Before making a choice, he can reduce the uncertainty in his prior belief μ by acquiring a costly signal about the filtered state \mathbf{Y}_t . Bob is restricted to signals and hence actions that are independent of \mathbf{X}_t conditionally on \mathbf{Y}_t . Given θ , \mathbf{A}_t , \mathbf{F}_t , and μ , Bob chooses a state-dependent distribution over actions

$$P_t(\cdot | \mathbf{Y}_t; \theta, \mathbf{A}_t, \mathbf{F}_t, \mu) \in \Delta(A(\mathbf{A}_t))$$

which is equivalent to choosing a costly signal and making a choice contingent on its realization (e.g. Matejka and McKay, 2015, Corollary 1) and hence standard in the rational inattention literature.⁴ The choice of P_t comes at a cost

$$K_{\theta}(P; \mathbf{A}, \mathbf{F}, \mu) = \kappa I_P(a; vec(\mathbf{A}\mathbf{X}\mathbf{F}^T))$$

that is linear in the Shannon mutual information $I_P(a; vec(\mathbf{Y}))$ under P between a and the vectorized filtered state $vec(\mathbf{Y})$. The problem of a receiver of type θ is thus, given \mathbf{A}_t , \mathbf{F}_t , and μ , to choose

$$P_t^* = \arg\max_{P_t} \int_{\mathcal{X}} \sum_{a \in A(\mathbf{A})} P_t(a | \mathbf{A}_t \mathbf{X}_t \mathbf{F}_t^T; \theta, \mathbf{A}_t, \mathbf{F}_t, \mu) u_{\theta}(a, \mathbf{X}_t) d\mu(\mathbf{X}_t) - K_{\theta}(P_t; \mathbf{A}_t, \mathbf{F}_t, \mu)$$
(3)

Alice's Information and Preferences Alice observes the state \mathbf{X}_t before choosing \mathbf{A}_t and \mathbf{F}_t , but not Bob's type θ . Instead, she only knows that the static type is distributed according to $\tau \in \Delta(\Theta)$. Alice's utility function $v(a, \mathbf{AX}, \theta)$ may nonetheless depend on Bob's type. Alice discounts Bob's information costs with a discount factor $\alpha \in [0, 1]$, so that her per-period utility is given by

$$R^{v}_{\alpha}(a, \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) = v(a, \mathbf{A}\mathbf{X}, \theta) - (1 - \alpha)K_{\theta}(P^{*}; \mathbf{A}, \mathbf{F}, \mu).$$
(4)

It will also be convenient to define the expected utility, conditional on **X** and θ , of Bob's stochastic choice function P^* under utility v and with discount factor $\alpha \in [0, 1]$ on Bob's

⁴Note in particular that Bob conditions his actions on \mathbf{Y}_t without (fully) observing it.



Figure 1: Illustration of the timeline within a period: The joint state $s_t := (\mathbf{X}_t, \theta)$ is realized at the beginning of the period. The Sender (Alice) observes \mathbf{X}_t but not θ , on which she maintains a prior belief $b_t \in \Delta(\Theta)$. Given \mathbf{X}_t and b_t , Alice chooses an information filter $(\mathbf{A}_t, \mathbf{F}_t)$. Bob observes $(\mathbf{A}_t, \mathbf{F}_t)$ as well as θ , but not \mathbf{X}_t , on which he maintains a prior belief $\mu \in \Delta(\mathcal{X})$. Bob's action a_t is realized from the distribution P_t^* he chooses given $\theta, \mathbf{A}_t, \mathbf{F}_t$, and μ . Upon observing a_t , Alice forms posterior belief b_{t+1} ; Bob's prior belief μ is stationary.

information costs,

$$U^{v}_{\alpha}(\mathbf{X},\theta,\mathbf{A},\mathbf{F};\mu) := \sum_{a \in \mathcal{A}} P^{*}(a|\mathbf{A}\mathbf{X}\mathbf{F}^{T};\theta,\mathbf{A},\mathbf{F},\mu) R^{v}_{\alpha}(a,\mathbf{X},\theta,\mathbf{A},\mathbf{F};\mu).$$
(5)

Timing The timing in each period $t \in \{0, ..., \infty\}$ is as follows:

- 1. Alice observes the realization of $\mathbf{X}_t \sim \mu$, but not θ (instead maintaining a prior belief $b_t \in \Delta(\Theta)$, with $b_0 = \tau$).
- 2. Given \mathbf{X}_t and b_t , Alice chooses \mathbf{A}_t and \mathbf{F}_t .
- 3. Bob observes θ , \mathbf{A}_t , and \mathbf{F}_t , but not \mathbf{X}_t (instead maintaining a prior belief $\mu \in \Delta(\mathcal{X})$).
- 4. Bob chooses $P_t^*(\cdot | \mathbf{Y}_t; \theta, \mathbf{A}_t, \mathbf{F}_t, \mu)$.
- 5. $a_t \sim P_t^*(\cdot | \mathbf{Y}_t; \theta, \mathbf{A}_t, \mathbf{F}_t, \mu)$ is realized. Bob receives utility $R_0^u(a_t, \mathbf{X}_t, \theta, \mathbf{A}_t, \mathbf{F}_t; \mu) = u_\theta(a_t, \mathbf{A}\mathbf{X}_t) - K_\theta(P^*; \mathbf{A}_t, \mathbf{F}_t, \mu)$. Alice observes a_t , forms posterior belief $b_{t+1}(\theta | a_0, \dots, a_t)$, and receives utility

$$\int_{\Theta} b_{t+1}(\theta | a_0, \dots, a_t) R^v_{\alpha}(a_t, \mathbf{X}_t, \theta, \mathbf{A}_t, \mathbf{F}_t; \mu) d\theta$$

3 Example

This section illustrates the model's solution for the special case of two actions (\tilde{a} and b) and two attributes (1 and 2). To build intuition we further assume (1) a single time period, (2) aligned utilities (v = u), (3) that θ is known to Alice (i.e., τ is degenerate), and (4) that $\mathbf{A} = \mathbf{I}_2$ is the identity matrix so that both actions are available to Bob, whose utility is given by $u_{\theta}(a, \mathbf{X}) = w_1 x_{\tilde{a}1} + w_2 x_{\tilde{a}2}$. In this section, we assume that the prior μ is a matrix Gaussian distribution

$$\mathbf{X} \sim \mathcal{MN}_{2 \times 2} \left(\mathbf{M} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, \mathbf{U} = \begin{bmatrix} \nu & c \\ c & \nu \end{bmatrix}, \mathbf{V} = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \right).$$

where **M** is the mean, **U** the action-covariance matrix, and **V** the attribute-covariance matrix, with ν the variance of both actions, c the covariance across actions, and ρ the covariance across attributes. Most results in this section should extend, for the case of two actions, to more general priors, but the matrix Gaussian distribution's parametrization will facilitate interpretation. Alice chooses among the following three information designs

$$\mathbf{F}^{1} = \begin{bmatrix} 1 & 0 \end{bmatrix} \qquad \mathbf{F}^{2} = \begin{bmatrix} 0 & 1 \end{bmatrix} \qquad \mathbf{F}^{12} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

which correspond to revealing the first attribute only, the second attribute only, and revealing both attributes, respectively. In order to study Alice's problem, we first have to understand Bob's solution.

Bob's Choice Behavior Bob's optimal conditional choice probabilities, as we will show in section 4, can be obtained as

$$P^{*}\left(\tilde{a} \mid \mathbf{Y}; \theta, \mathbf{F}^{1}, \mu\right) = \left[1 + \exp\left(-\frac{w_{1} + \rho w_{2}}{\kappa}(Y_{\tilde{a}1} - Y_{\tilde{b}1})\right)\right]^{-1}$$

$$P^{*}\left(\tilde{a} \mid \mathbf{Y}; \theta, \mathbf{F}^{2}, \mu\right) = \left[1 + \exp\left(-\frac{\rho w_{1} + w_{2}}{\kappa}(Y_{\tilde{a}2} - Y_{\tilde{b}2})\right)\right]^{-1}$$

$$P^{*}\left(\tilde{a} \mid \mathbf{Y}; \theta, \mathbf{F}^{12}, \mu\right) = \left[1 + \exp\left(-\frac{w_{1}}{\kappa}(Y_{\tilde{a}1} - Y_{\tilde{b}1}) - \frac{w_{2}}{\kappa}(Y_{\tilde{a}2} - Y_{\tilde{b}2})\right)\right]^{-1}.$$

This solution parsimoniously captures a number of intuitive characteristics. First, Bob's probability of choosing action \tilde{a} depends on how it differs from action \tilde{b} in terms of the attributes that are accessible to him, weighted by his preference weights. Second, Bob makes an inference based on the observable attribute about any potentially unobservable

ones, relying on the attribute covariance ρ . Third, the accuracy of his choices depends on the ratio between the utility difference across actions and the information costs κ .

Bob's log-odds of choosing \tilde{a} are a function of $\boldsymbol{dX} = \begin{bmatrix} dX_1 \\ dX_2 \end{bmatrix} = \begin{bmatrix} X_{\tilde{a}1} - X_{\tilde{b}1} \\ X_{\tilde{a}2} - X_{\tilde{b}2} \end{bmatrix}$, given by

$$\Delta_{\mathbf{F}}(\boldsymbol{dX};\boldsymbol{\theta},\boldsymbol{\mu}) \equiv \ln \frac{P^*\left(\tilde{a} \left| \mathbf{X}\mathbf{F}^T;\boldsymbol{\theta},\mathbf{F},\boldsymbol{\mu} \right. \right)}{P^*\left(\tilde{b} \left| \mathbf{X}\mathbf{F}^T;\boldsymbol{\theta},\mathbf{F},\boldsymbol{\mu} \right. \right)} = \begin{cases} \frac{w_1 + \rho w_2}{\kappa} dX_1 & \text{if } \mathbf{F} = \mathbf{F}^1 \\ \frac{\rho w_1 + w_2}{\kappa} dX_2 & \text{if } \mathbf{F} = \mathbf{F}^2 \\ \frac{w_1}{\kappa} dX_1 + \frac{w_2}{\kappa} dX_2 & \text{if } \mathbf{F} = \mathbf{F}^{12} \end{cases}$$

Welfare-Maximizing Information Policy These log-odds are also a sufficient statistic for maximizing Bob's gross welfare

$$U_1^u(\mathbf{X},\theta,\mathbf{F};\mu) = \frac{w_1(X_{\tilde{a}1} - X_{\tilde{b}1}) + w_2(X_{\tilde{a}2} - X_{\tilde{b}2})}{1 + \exp\left(-\Delta_{\mathbf{F}}(d\mathbf{X})\right)} + w_1X_{\tilde{b}1} + w_2X_{\tilde{b}2},$$

which is increasing in $\Delta_{\mathbf{F}}(d\mathbf{X})$ if and only if $\mathbf{w}^T d\mathbf{X} \ge 0$, so that

$$U_{1}^{u}\left(\mathbf{X}, \theta, \mathbf{F}^{1}; \mu\right) \geq U_{1}^{u}\left(\mathbf{X}, \theta, \mathbf{F}^{12}; \mu\right) \iff \operatorname{sgn}(\boldsymbol{w}^{T}\boldsymbol{d}\boldsymbol{X}) \begin{bmatrix} \rho \\ -1 \end{bmatrix}^{T} \boldsymbol{d}\boldsymbol{X} \geq 0$$
$$U_{1}^{u}\left(\mathbf{X}, \theta, \mathbf{F}^{2}; \mu\right) \geq U_{1}^{u}\left(\mathbf{X}, \theta, \mathbf{F}^{12}; \mu\right) \iff \operatorname{sgn}(\boldsymbol{w}^{T}\boldsymbol{d}\boldsymbol{X}) \begin{bmatrix} -1 \\ \rho \end{bmatrix}^{T} \boldsymbol{d}\boldsymbol{X} \geq 0$$
$$U_{1}^{u}\left(\mathbf{X}, \theta, \mathbf{F}^{1}; \mu\right) \geq U_{1}^{u}\left(\mathbf{X}, \theta, \mathbf{F}^{2}; \mu\right) \iff \operatorname{sgn}(\boldsymbol{w}^{T}\boldsymbol{d}\boldsymbol{X}) \begin{bmatrix} w_{1} + \rho w_{2} \\ -(\rho w_{1} + w_{2}) \end{bmatrix}^{T} \boldsymbol{d}\boldsymbol{X} \geq 0$$

The gross-welfare-maximizing information policy \mathbf{F} is thus determined by three hyperplanes that partition the space of possible $d\mathbf{X}$, in a manner that depends on Bob's preference weights \mathbf{w} and the attribute covariance ρ under his prior belief. Figure 2 shows the gross welfare under the three information policies as a function of the state, for a number of different parameters. The last column illustrates how the hyperplanes determine the optimal policy as a function of the underlying parameters. Figure 7 in the appendix shows the same information as a function of preference weights \mathbf{w} , for a number of different state realizations. Note that the gross-welfare-maximizing information policy is independent of the scale of preference weights; only their sign and ratio matters. This also means that the level of information costs $\kappa > 0$ is irrelevant for the gross-welfare-maximizing policy.

Interpreting the Welfare-Maximizing Information Policy One feature of the grosswelfare-maximizing information policy is particularly salient: The set of states for which



Figure 2: The plots in the first three columns show the gross welfare $U_1^u(\mathbf{X}, \theta, \mathbf{F}; \mu)$ for the three different information policies \mathbf{F} , as a function of $X_{\tilde{a}1} - X_{\tilde{b}1}$ and $X_{\tilde{a}2} - X_{\tilde{b}2}$. The last column shows which of the information policies maximize gross welfare for any given \mathbf{X} , with the color bar associating colors with the three information policies. Each row corresponds to a set of parameters including Bob's preference type and the attribute correlation ρ under his prior belief ($\nu = 1$ and c = 0 in all rows).

Choice	Attribute 1	Attribute 2	w_1	w_2	ρ
coffee shops	taste	friendliness	1	1	0
hotels	rating	beach proximity	1	1	0.5
flights	duration	price	-1	-1	-0.5
wines	quality	price	1	-1	0.75

Table 1: Summary of the examples described in the text, as illustrated in Figure 2.

occluding information is optimal is determined by the attribute covariance ρ . In order to better understand the optimal policy, we will illustrate each row of Figure 2 with a brief example.

Uncorrelated, Positive Attributes First, consider a customer Bob choosing between two coffee shops whose coffee's taste $(w_1 = 1)$ and barista's friendliness $(w_2 = 1)$, which for the sake of this example are uncorrelated, he values equally. This is the case covered by the top row of Figure 2. If both attributes are better at one coffee shop, then the welfaremaximizing information policy gives Bob access to both attributes. If the attributes are conflicting, however, then the optimal information policy only gives Bob access to the more "decisive" attribute for which the difference across coffee shops is larger. The intuition is that Bob attends and is responsive to all available attributes, which improves his decision when both attributes align, but harms it when they are in conflict.

Positively Correlated, Positive Attributes The second row can be interpreted as modelling the choice between two hotels by a consumer who values both their rating $(w_1 = 1)$ and their proximity to the beach $(w_2 = 1)$, which happen to be positively correlated. Here, Alice should only make the more decisive attribute accessible, unless both attributes favor one hotel by a similar margin. In a scenario in which one hotel has a somewhat larger rating than the other but is also significantly closer to the beach, occluding the hotel ratings increases the hotel guest's welfare (in contrast to the coffee consumer). The reason is that the hotel guest will attend more to beach proximity in absence of a hotel rating than in its presence, because the correlation allows them to substitute acquiring information on beach proximity for information on hotel ratings. In states in which beach proximity differs by a sufficiently larger margin, this will increase their overall welfare.

Negatively Correlated, Negative Attributes Next, consider an airline customer choosing between two flights who incurs disutility from their duration $(w_1 = -1)$ as well as their price $(w_2 = -1)$, which are negatively correlated. Here, the Receiver's utility is greatest when both attributes are accessible, unless the attributes favour different flights but by a

similar margin. In a scenario in which one flight is somewhat shorter but significantly pricier, occluding the duration would harm the airline customer (in contrast to the coffee consumer and hotel guest). The reason is that the absence of duration information decreases the benefit of attending to the price – even though they are correlated – because the difference in expected utility across flights is small conditional on the price (while the difference in expected utility across hotels is large conditional on beach proximity).

Positively Correlated, Opposite Attributes A similar intuition holds in the scenario of Alice selling wine to Bob whose utility is increasing in quality $(w_1 = 1)$ but decreasing in price $(w_2 = -1)$, which are positively correlated. A (benevolent) Alice should make the wine quality accessible to Bob by way of a tasting, unless the difference in the wines' quality is very close to their price difference. Again, occluding wine quality will harm Bob even if a wine of moderately higher quality is also much more expensive, for example.

Elasticity of Demand with Respect to Attribute Values This intuition is also reflected in the elasticity of the probability of choosing option a with respect to the value of its attribute f, which is given by

$$\frac{\partial P^*(a|\mathbf{X}\mathbf{F}^T;\theta,\mathbf{F},\mu)}{\partial X_{af}}\frac{X_{af}}{P^*(a|\mathbf{X}\mathbf{F}^T;\theta,\mathbf{F},\mu)} = \frac{\exp(-\Delta_{\mathbf{F}}(\mathbf{d}\mathbf{X}))}{1+\exp(-\Delta_{\mathbf{F}}(\mathbf{d}\mathbf{X}))}\frac{\partial \Delta_{\mathbf{F}}(\mathbf{d}\mathbf{X})}{\partial X_{af}}X_{af}$$

Bob's price-elasticity of demand at $dX_1 = dX_2 = 0$ in the presence of information on quality $\left(\frac{w_2}{2\kappa}X_{\tilde{a}2}\right)$ is more negative (larger in magnitude) than in its absence $\left(\frac{\rho w_1 + w_2}{2\kappa}X_{\tilde{a}2}\right)$ if and only if $\rho w_1 > 0$, as is indeed the case in the wine example. The hotel guest's demand elasticity with respect to beach proximity is less positive (smaller in magnitude) in the presence of a hotel rating than in its absence.

The difference between these scenarios is a feature inherent to endogenous information acquisition that would not arise under exogenous signals. The wine consumer's utility difference across wines, $w_1 dX_1 + w_2 dX_2$, is more likely to be small than the hotel guest's. This implies that the wine (and airline!) customers have an overall lower incentive to acquire information than the hotel guest, even when the absolute correlation is the same. This becomes apparent from the distribution of $\Delta_{\mathbf{F}}(\mathbf{dX})$ induced by the prior $\mu(\mathbf{X})$,

$$\begin{split} \Delta_{\mathbf{F}^{1}}(\mathbf{d}\mathbf{X}) &\sim \mathcal{N}\left(0, 2(\nu - c)\left(\frac{w_{1} + \rho w_{2}}{\kappa}\right)^{2}\right)\\ \Delta_{\mathbf{F}^{2}}(\mathbf{d}\mathbf{X}) &\sim \mathcal{N}\left(0, 2(\nu - c)\left(\frac{\rho w_{1} + w_{2}}{\kappa}\right)^{2}\right)\\ \Delta_{\mathbf{F}^{12}}(\mathbf{d}\mathbf{X}) &\sim \mathcal{N}\left(0, 2(\nu - c)\frac{w_{1}^{2} + 2\rho w_{1}w_{2} + w_{2}^{2}}{\kappa^{2}}\right) \end{split}$$

The distribution of $\Delta_{\mathbf{F}}$ turns out to be a sufficient statistic for the amount of information Bob acquires. In the wine example, the variance of $\Delta_{\mathbf{F}^2}$ captures Bob's incentive to learn about the price attribute in absence of information on quality. The incentive to learn about the price in the presence of information on quality is reflected correspondingly in the variance of the distribution of $\Delta_{\mathbf{F}^{12}}$ conditional on dX_1 ,

$$\Delta_{\mathbf{F}^{12}}(\mathbf{dX}) | dX_1 \sim \mathcal{N}\left(\left(\frac{w_1}{\kappa} + \rho \frac{w_2}{\kappa}\right) dX_1, 2(\nu - c)(1 - \rho^2) \left(\frac{w_2}{\kappa}\right)^2\right).$$

The incentive to attend to the price is thus larger in presence of quality information than in its absence whenever

$$2\rho w_1 w_2 < -\rho^2 (w_1^2 + w_2^2)$$

which is indeed the case in the wine example, but not in the hotel example. This explains the difference in interaction: Conditioning on hotel ratings would decrease the dispersion in Δ and thus lower his learning incentive, so that acquiring information on ratings is a substitute for acquiring information on beach proximity. Conditioning on wine quality, on the other hand, increases the dispersion in Δ and hence raises learning incentives, so that acquiring information on the price.

Information Costs and Net Welfare The distribution of $\Delta_{\mathbf{F}}$ is thus also a sufficient statistic for the information costs of the optimal strategy, which are given by

$$K_{\theta}\left(P^{*};\mathbf{F},\mu\right) = \kappa \mathbb{E}_{\Delta_{\mathbf{F}}}\left[\frac{\Delta_{\mathbf{F}}\exp(\Delta_{\mathbf{F}})}{1+\exp(\Delta_{\mathbf{F}})} + \ln\left(\frac{2}{1+\exp(\Delta_{\mathbf{F}})}\right)\right].$$

The term in the expectation is the Kullback-Leibler divergence from a Bernoulli random variable with log-odds 0 to a Bernoulli random variable with log-odds $\Delta_{\mathbf{F}}$, which is symmetric around 0 and strictly increasing in $|\Delta_{\mathbf{F}}|$. Given the centered Gaussian distribution of $\Delta_{\mathbf{F}}$, the information costs are thus increasing in (and pinned down by) the variance $\operatorname{Var}(\Delta_{\mathbf{F}})$.



Figure 3: The plots in the first three columns show the net welfare $U_0^u(\theta, \mathbf{X}, \mathbf{F}; \mu)$ for the three different information policies \mathbf{F} , as a function of $X_{\tilde{a}1} - X_{\tilde{b}1}$ and $X_{\tilde{a}2} - X_{\tilde{b}2}$. The last column shows which of the information policies maximize gross welfare for any given \mathbf{X} , with the color bar associating colors with information policies. Each row corresponds to a set of parameters.

The resulting expected utility for a Sender internalizing a fraction $(1 - \alpha)$ of information costs is given by

$$\begin{aligned} U^{u}_{\alpha}\left(\mathbf{X}, \theta, \mathbf{F}; \mu\right) &= \\ \frac{\Delta_{\mathbf{F}}(\boldsymbol{d}\mathbf{X})}{1 + \exp\left(-\Delta_{\mathbf{F}}(\boldsymbol{d}\mathbf{X})\right)} - (1 - \alpha)\kappa \mathbb{E}_{\Delta_{\mathbf{F}}}\left[\frac{\Delta_{\mathbf{F}}}{1 + \exp(-\Delta_{\mathbf{F}})} + \ln\left(\frac{2}{1 + \exp(\Delta_{\mathbf{F}})}\right)\right] + w_{1}X_{\tilde{b}1} + w_{2}X_{\tilde{b}2}. \end{aligned}$$

Figure 3 illustrates the net utility ($\alpha = 0$) along with the net-welfare-maximizing information policy: Alice should only give Bob access to the one attribute that is more decisive, unless information costs are sufficiently low (high $|\boldsymbol{w}|/\kappa$) that having access to both attributes is better. The net-welfare-maximizing information policy thus does depend on the level of \boldsymbol{w}/κ , in contrast to the gross-welfare-maximizing policy.

4 Receiver Problem

This section discusses the solution to Bob's (static) problem

$$\overline{U}_{0}^{u}(\theta, \mathbf{A}, \mathbf{F}; \mu) = \max_{P} \mathbb{E}_{\mathbf{X} \sim \mu} \left[\sum_{a \in A(\mathbf{A})} P(a | \mathbf{A} \mathbf{X} \mathbf{F}^{T}; \theta, \mathbf{A}, \mathbf{F}, \mu) R_{0}^{u}(a, \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) \right]$$

We will first consider how Bob updates his prior belief when observing the realization of a signal on *part of* the state. We will then provide a solution to Bob's problem in the general case. Next, we will provide some intuition by discussing the solution in the special case of a Gaussian prior belief. Lastly, we will provide a formula for the welfare resulting under Bob's optimal strategy.

4.1 Conditional Beliefs

As outlined in section 2, Bob can acquire costly information on the filtered state \mathbf{Y} only. The set of states consistent with a given filtered state \mathbf{Y} is defined as

$$\chi(\mathbf{Y}; \mathbf{A}, \mathbf{F}) \equiv \left\{ \mathbf{X} \in \mathcal{X} : \mathbf{A}\mathbf{X}\mathbf{F}^T = \mathbf{Y} \right\}.$$

A filter (\mathbf{A}, \mathbf{F}) is thus a deterministic signal structure pooling states, i.e. partitioning \mathcal{X} into $\{\chi(\mathbf{Y}; \mathbf{A}, \mathbf{F})\}_{\mathbf{Y}}$. The sets of pooled states are characterized by the following fact from linear algebra.

Lemma 1. Let $\mathbf{H} = \mathbf{A} \otimes \mathbf{F}$ and $\mathbf{H}^+ = \mathbf{H}^T (\mathbf{H} \mathbf{H}^T)^{-1}$ its (right) pseudo-inverse. Then

$$\mathbf{X} \in \chi(\mathbf{Y}; \mathbf{A}, \mathbf{F}) \Leftrightarrow vec(\mathbf{X}^T) \in \left\{ \mathbf{H}^+ vec(\mathbf{Y}^T) + (\mathbf{I}_{m \cdot n} - \mathbf{H}^+ \mathbf{H}) \boldsymbol{z} : \boldsymbol{z} \in \mathbb{R}^{m \cdot n} \right\}.$$

Conditional on a filtered state \mathbf{Y} , Bob's conditional belief over the set of states \mathcal{X} will then be given by

$$\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F}) = \begin{cases} \frac{\mu(\mathbf{X})}{\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F})} & \text{if } \mathbf{X} \in \chi(\mathbf{Y}; \mathbf{A}, \mathbf{F}) \\ 0 & \text{otherwise} \end{cases}$$
(6)

where

$$\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) \equiv \int_{\chi(\mathbf{Y}; \mathbf{A}, \mathbf{F})} d\mu(\mathbf{X}).$$
(7)

In deciding what information to acquire, Bob will consider this conditional updating in evaluating the benefit of acquiring information. In solving Bob's problem, it will be helpful that the prior's row-exchangeability is preserved under any permissible information filter, as is established by the following Lemma.

Lemma 2. Given any $k \times k$ permutation matrix \mathbf{P}_k , let $\mathbf{P}_m^{\mathbf{P}_k} := \mathbf{A}^+ \mathbf{P}_k \mathbf{A} + \mathbf{I}_m - \mathbf{A}^+ \mathbf{A}$ where $\mathbf{A}^+ := \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1}$. If $\mu(\mathbf{X})$ is row-exchangeable, then for any $k \times k$ permutation matrix \mathbf{P}_k and any (\mathbf{A}, \mathbf{F}) , it is the case that

$$\mu_{\mathbf{Y}}(\mathbf{P}_{k}\mathbf{Y}|\mathbf{A},\mathbf{F}) = \mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A},\mathbf{F}) \qquad \forall \mathbf{Y} \in \mathcal{Y}$$
(8)

and

$$\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X}|\mathbf{P}_{k}\mathbf{Y},\mathbf{A},\mathbf{F}) = \mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y},\mathbf{A},\mathbf{F}) \qquad \forall \mathbf{Y} \in \mathcal{Y}, \mathbf{X} \in \chi(\mathbf{Y};\mathbf{A},\mathbf{F}).$$
(9)

4.2 Optimal Stochastic Choice Strategy

With these results in place, we can now state the solution to Bob's problem.

Theorem 1. Given a row-exchangeable prior μ , the receiver problem (eq. 3) has a solution

$$P^*(a|\mathbf{Y};\theta,\mathbf{A},\mathbf{F},\mu) = \frac{z_{\theta}(a,\mathbf{Y};\mathbf{A},\mathbf{F},\mu)}{\sum_{c\in A(\mathbf{A})} z_{\theta}(c,\mathbf{Y};\mathbf{A},\mathbf{F},\mu)} \qquad \forall a \in A(\mathbf{A})$$
(10)

and $P^*(a|\mathbf{Y}; \theta, \mathbf{A}, \mathbf{F}, \mu) = 0$ for all $a \notin A(\mathbf{A})$, where

$$z_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) = \exp\left(\frac{1}{\kappa} \int_{\mathcal{X}} u_{\theta}(a, \mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F})\right).$$

This theorem states that Bob's problem with the row-exchangeable prior of Definition 1 is solved by conditional choice probabilities of the (unbiased) multinomial logit form. According to this solution, the probability of choosing an action is proportional to its expected utility relative to the marginal cost of information κ .

This result provides an analytical solution to a rational inattention problem that is nonstandard in that it features a partially observable matrix state. It generalizes the result of Bucher and Caplin (2021) to identify a condition on the prior belief (which here is a matrix distribution) under which conditional choice probabilities of the unbiased multinomial logit form are not just necessary (Matejka and McKay, 2015), but also satisfy the sufficient conditions for a solution (Caplin et al., 2019). For general prior beliefs, the unconditional choice probabilities in the solution of Matejka and McKay (2015) have to be pinned down numerically using the Blahut-Arimoto algorithm, and corner solutions may give rise to endogenous consideration sets (Caplin et al., 2019). The row-exchangeable prior beliefs of Definition 1 guarantee that uniform unconditional choice probabilities are optimal.

The fact that Bob can acquire information only on a subset of the state has implications for his information acquisition, because his incentive to acquire information depends on the conditional distribution $\mu_{\mathbf{X}|\mathbf{Y}}$. The example of section 3 has demonstrated, for instance, that correlated attributes may increase the incentive to acquire information on observable attributes in order to infer something about hidden ones.

4.3 Special Case: Matrix Gaussian Prior and Linear Preferences

This subsection discusses a special case to build better intuition for the solution of Theorem 1: We assume that Bob's utility is linear in attributes, with preference weights $\boldsymbol{w} \in \mathbb{R}^n$,

$$u_{\theta}(a, \mathbf{A}\mathbf{X}) = \boldsymbol{e}_{a}^{T}\mathbf{A}\mathbf{X}\boldsymbol{w} \qquad \forall a \in A(\mathbf{A}),$$
(11)

where e_a is the k-dimensional standard basis vector representing the chosen action $a \in A(\mathbf{A})$. Lemma 3. The linear utility function $u_{\theta}(a, \mathbf{AX}) = e_a^T \mathbf{AX} \boldsymbol{w}$ is invariant under permutation.

We further assume that Bob's prior belief μ about **X** is a matrix Gaussian distribution

$$\mathbf{X} \sim \mathcal{MN}_{m \times n}(\mathbf{M}, \mathbf{U}, \mathbf{V}) \tag{12}$$

with a mean $\mathbf{M} = \mathbf{1}_m \otimes \mathbf{m}^T \in \mathbb{R}^{m \times n}$ consisting of identical rows $\mathbf{m} \in \mathbb{R}^n$, a so-called "completely symmetric" action covariance matrix $\mathbf{U} = c\mathbf{1}_m\mathbf{1}_m^T + (\nu - c)\mathbf{I}_m \in \mathbb{R}^{m \times m}$, and an arbitrary attribute covariance matrix $\mathbf{V} \in \mathbb{R}^{n \times n}$. These assumptions are sufficient for the matrix Gaussian distribution to be row-exchangeable.

Lemma 4. The matrix Gaussian distribution of equation 12 is row-exchangeable (Def. 1) if $\mathbf{M} = \mathbf{1}_m \otimes \mathbf{m}^T$ and $\mathbf{U} = c\mathbf{1}_m\mathbf{1}_m^T + (\nu - c)\mathbf{I}_m$.

Given the matrix Gaussian prior, the (marginal) distribution $\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F})$ of the filtered state is given by

$$\mathbf{Y} \equiv \mathbf{A}\mathbf{X}\mathbf{F}^T \sim \mathcal{M}\mathcal{N}_{k \times l}(\mathbf{A}\mathbf{M}\mathbf{F}^T, \mathbf{A}\mathbf{U}\mathbf{A}^T, \mathbf{F}\mathbf{V}\mathbf{F}^T).$$

With linear preferences, only the mean of the conditional distribution $\mu_{\mathbf{X}|\mathbf{Y}}$ will appear in the solution to be stated in the Corollary below. To state the conditional mean, it will be helpful to define the matrix $\overline{\mathbf{F}}$ as the "complement" of \mathbf{F} , in the sense that it consists of the non-zero rows of matrix $\mathbf{I}_n - diag(\mathbf{1}_n^T \mathbf{F})$ (if any; otherwise $\overline{\mathbf{F}} = 0$) and hence collects all attributes that are not accessible under \mathbf{F} .

Corollary 1. Given the matrix Gaussian prior of equation 12, the receiver problem (eq. 3) has a solution

$$P(a|\mathbf{Y};\theta,\mathbf{A},\mathbf{F},\mu) = \frac{\exp\left(\frac{1}{\kappa}\boldsymbol{e}_{a}^{T}\mathbf{A}\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}\left[\mathbf{X}|\mathbf{Y},\mathbf{A},\mathbf{F}\right]\boldsymbol{w}\right)}{\sum_{c\in A(\mathbf{A})}\exp\left(\frac{1}{\kappa}\boldsymbol{e}_{c}^{T}\mathbf{A}\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}\left[\mathbf{X}|\mathbf{Y},\mathbf{A},\mathbf{F}\right]\boldsymbol{w}\right)} \qquad \forall a \in A(\mathbf{A})$$

where

$$\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}\left[\mathbf{A}\mathbf{X}|\mathbf{Y},\mathbf{A},\mathbf{F}\right] = \mathbf{Y}\mathbf{F} + \left(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T} + (\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^{T}\right)\overline{\mathbf{F}}.$$

Note that with linear preferences, the probability of choosing an action notably depends on the conditional mean of its attribute values under the prior belief, weighted by the preference weights \boldsymbol{w} . The mean conditions on the observable part \mathbf{Y} of the state \mathbf{X} , as determined by the filter (\mathbf{A}, \mathbf{F}) . To interpret the conditional expectation, note that the left term $\mathbf{YF} \in \mathbb{R}^{k \times n}$ simply projects the observable state \mathbf{Y} back into the original attribute space. The hidden attributes will be missing, but can be imputed by the term in parentheses as follows: Their unconditional expectation is given by the first summand \mathbf{AMF}^T . The second summand in parentheses captures what the observable attributes imply about the unobservable ones: The deviation of \mathbf{Y} from its expectation \mathbf{AMF}^T allows an inference about how the unobservable attributes deviate from their expectation \mathbf{AMF}^T , to an extent that depends on the "ratio" of the covariances of unobservable and observable attributes. Both summands are then projected back to the original attribute space by right-multiplying with $\mathbf{\overline{F}}$ before adding them to the observable attributes \mathbf{YF} .

Note in particular that the inference and hence solution depend on the attribute covariance \mathbf{V} , but *not* on the action covariance \mathbf{U} . This implies that Bob's belief about the correlation among the exchangeable choice options is irrelevant to Alice. The solution's dependence on the filter (\mathbf{A}, \mathbf{F}) makes explicit how Alice is able to affect Bob's choices.

4.4 Welfare

Based on Bob's optimal stochastic choice strategy as given by Theorem 1, we can obtain its resulting expected welfare.

Proposition 1. The expected welfare, as defined in equation 5, of Bob's optimal stochastic choice strategy P^* , evaluated under a utility function v, is given by

$$U_{\alpha}^{v}(\mathbf{X},\theta,\mathbf{A},\mathbf{F};\mu) = \sum_{a\in A(\mathbf{A})} \left[\frac{z_{\theta}(a,\mathbf{A}\mathbf{X}\mathbf{F}^{T})v(a,\mathbf{A}\mathbf{X},\theta)}{\sum_{c\in A(\mathbf{A})} z_{\theta}(c,\mathbf{A}\mathbf{X}\mathbf{F}^{T})} - (1-\alpha) \int_{\mathbf{Y}\in\mathcal{Y}} \frac{z_{\theta}(a,\mathbf{Y})\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}\left[u_{\theta}(a,\mathbf{A}\mathbf{X})|\mathbf{Y}\right]}{\sum_{c\in A(\mathbf{A})} z_{\theta}(c,\mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A},\mathbf{F}) \right]$$
(13)

$$+ (1 - \alpha)\kappa \int_{\mathbf{Y} \in \mathcal{Y}} \ln\left(\frac{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})}{|A(\mathbf{A})|}\right) d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F})$$
(14)

where

$$\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}\left[u_{\theta}(a, \mathbf{A}\mathbf{X})|\mathbf{Y}, \mathbf{A}, \mathbf{F}\right] = \int_{\mathbf{X} \in \mathcal{X}} u_{\theta}(a, \mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F})$$

and

$$z_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) = \exp\left(\frac{1}{\kappa} \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}\left[u_{\theta}(a, \mathbf{A}\mathbf{X})|\mathbf{Y}, \mathbf{A}, \mathbf{F}\right]\right)$$

This result explicitly states the welfare resulting from Bob's solution to his rational inattention problem. Setting $\alpha = 1$ yields the gross welfare

$$U_1^v(\mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) = \sum_{a \in A(\mathbf{A})} \frac{z_{\theta}(a, \mathbf{A}\mathbf{X}\mathbf{F}^T)v(a, \mathbf{A}\mathbf{X}, \theta)}{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{A}\mathbf{X}\mathbf{F}^T)},$$

which takes a particularly elementary form, considering the intricacies of Bob's problem. On the other end of the spectrum, Bob's net welfare ($\alpha = 0$) – for the case in which $rank(\mathbf{A}) = m$ and $rank(\mathbf{F}) = n$ so that $\mu_{\mathbf{X}|\mathbf{Y}}$ is degenerate – reduces to

$$U_0^u(\mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) = \kappa \ln \left(\frac{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{A}\mathbf{X}\mathbf{F}^T)}{|A(\mathbf{A})|} \right).$$

5 Sender Problem

In the example of section 3, Alice was assumed to observe Bob's type before choosing an information filter (\mathbf{A}, \mathbf{F}) . In this section, we now turn to solving Alice's general problem, as specified in section 2, where this is not the case. Alice's problem is dynamic, with a temporal discount factor γ . Recall that at the beginning of each period t, Alice observes \mathbf{X}_t but not θ_t . Upon choosing a filter $(\mathbf{A}_t, \mathbf{F}_t)$, Alice observes Bob's choice $a_t \sim P^*(\cdot | \mathbf{A} \mathbf{X} \mathbf{F}^T; \theta, \mathbf{A}, \mathbf{F}, \mu)$.

5.1 Partially Observable Markov Decision Process

Alice's problem can be formulated as a partially observable Markov decision process (POMDP; Kaelbling et al., 1998) defined as the tuple $\langle S, \mathcal{H}, \mathcal{T}, R^v_{\alpha}, \mathcal{A}, O, \gamma \rangle$ consisting of the (latent) state space $S = \mathcal{X} \times \Theta$, Alice's set of filter choices $\mathcal{H} = \mathcal{A} \times \mathcal{F}$, the state transition kernel

$$\mathcal{T}(\mathbf{X}', \theta' | \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}) = \mu(\mathbf{X}') \tau(\theta' | \theta),$$

where $\tau(\theta'|\theta) = \mathbf{1}_{\{\theta\}}(\theta')$, Alice's reward $R^v_{\alpha}(a, \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu)$, Alice's set of observations consisting of the set \mathcal{A} of Bob's choices, and their distribution

$$O(a|\mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) = P^*(a|\mathbf{A}\mathbf{X}\mathbf{F}^T; \theta, \mathbf{A}, \mathbf{F}, \mu),$$

as well as Alice's discount factor γ .⁵

Since Alice maintains a belief $b_t \in \Delta(\Theta)$, with $b_0 = \tau$, this gives rise to the following belief MDP (Markov decision process). Let $V_t(\mathbf{X}_t, b_t, \mu_t)$ be the value at the beginning of period t of the problem of the sender who has observed \mathbf{X}_t , holds beliefs $b_t \in \Delta(\Theta)$ and $\mu_t \in \Delta(\mathcal{X})$, and who is about to choose an information filter $\mathbf{A}_t, \mathbf{F}_t$. Alice's problem is thus to find a policy $\pi(\mathbf{A}, \mathbf{F} | \mathbf{X}, b, \mu)$ that satisfies the Bellman optimality equation

$$V^{\pi*}(\mathbf{X}, b, \mu) = \max_{\mathbf{A}, \mathbf{F}} \mathbb{E}_{\theta \sim b} \left[\sum_{a \in \mathcal{A}} P^*(a | \mathbf{A} \mathbf{X} \mathbf{F}^T; \theta, \mathbf{A}, \mathbf{F}, \mu) \left(R^v_{\alpha}(a, \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) + \gamma \mathbb{E}_{\mathbf{X}' \sim \mu} [V^{\pi*}(\mathbf{X}', b'(a), \mu)] \right) \right]$$
(15)

where the posterior belief b' is obtained as

$$b'(\theta'|a; \mathbf{X}, b, \mathbf{A}, \mathbf{F}, \mu) = \frac{\sum_{\theta \in \Theta} P^*(a | \mathbf{A} \mathbf{X} \mathbf{F}^T; \theta, \mathbf{A}, \mathbf{F}, \mu) \tau(\theta'|\theta) b(\theta)}{\sum_{\theta \in \Theta} P^*(a | \mathbf{A} \mathbf{X} \mathbf{F}^T; \theta, \mathbf{A}, \mathbf{F}, \mu) \sum_{\tilde{\theta} \in \Theta} \tau(\tilde{\theta}|\theta) b(\theta)}.$$
 (16)

Note that this Bayesian belief update is similar in spirit to the sharp revealed preference

⁵The transition kernel \mathcal{T} could in principle be unknown to Alice.



Figure 4: Cumulative mean regret ϱ_t^{π} (as defined in equation 17) under Alice's policy as well as the full-information policy $\mathbf{F} = \mathbf{I}$.

results of Caplin and Martin (2021) and Caplin et al. (2022) characterizing a set of Θ consistent with observed behavior, which could be interpreted as the support of $b'(\theta')$.

5.2 Numerical Solution

Solving the dynamic program of equation 15 is not possible analytically, so we find an approximate numerical solution using partially observable Monte Carlo planning (POMCP; Silver and Veness, 2010), the state-of-the-art algorithm for solving POMDPs. POMCP relies on Monte-Carlo tree search (MCTS) in combination with a particle filter to perform the belief update corresponding to equation 16. Our implementation relies on the POMDPs.jl package (Egorov et al., 2017).

5.3 Simulation Results

The resulting simulations demonstrate that Alice succeeds in learning an information policy that outperforms the benchmark of a naïve information policy. Here, we present the result of a simulation experiment making the same assumptions as the example of section 3, except that there are now T = 250 time periods, and Alice's initial belief b_0 is non-degenerate: We assume that b_0 is uniform over $\boldsymbol{w} = (\frac{2}{3}, \frac{1}{3})$ and $\boldsymbol{w} = (\frac{1}{3}, \frac{2}{3})$, while $\kappa = 1$ is known by Alice. We still consider aligned utilities (v = u) with Alice maximizing Bob's gross welfare ($\alpha = 1$). Focusing on Alice's choice of \mathbf{F} , we again assume that $\mathbf{A} = \mathbf{I}_2$. Figure 4 plots the cumulative



Figure 5: Evolution of Alice's belief $b_t(\theta)$ (top) and filter choices \mathbf{F}_t (bottom; moving average over 25 time periods). The bottom plot also displays the frequency of filter choices maximizing the instantaneous expected utility ("accuracy").

expected regret (cf. Loomes and Sugden, 1982)

$$\varrho_t^{\pi}(\{\mathbf{X}_{\tau}, \mathbf{A}_{\tau}, \mathbf{F}_{\tau}\}_{\tau=0}^t, \theta; \mu) = \sum_{\tau=0}^t \left[\max_{\mathbf{A}, \mathbf{F}} U_{\alpha}^{\upsilon}(\mathbf{X}_{\tau}, \theta, \mathbf{A}, \mathbf{F}; \mu) - U_{\alpha}^{\upsilon}(\mathbf{X}_{\tau}, \theta, \mathbf{A}_{\tau}, \mathbf{F}_{\tau}; \mu) \right]$$
(17)

as a function of time. It demonstrates how Alice's policy outperforms the full-information policy in that its regret is lower. Under Alice's policy, there is a point at which regret stops accumulating, indicating that Alice has succeeded in learning the optimal policy. Under the full-information policy, on the other hand, regret continues to accumulate.

In order to better understand how Alice learns the optimal policy, Figure 5 shows the corresponding evolution of Alice's belief $b_t(\theta)$ and the filter choices \mathbf{F}_t . Alice's initial belief is uniform over Bob's two possible types, but evolves as Alice observes Bob's actions, and eventually converges to Bob's true type. Alice's filter choices, shown in the bottom panel of Figure 5, respond both to the changing \mathbf{X}_t and the evolving belief b_t . The fraction of filter choices maximizing instantaneous expected utility reflects Alice's evolving belief: As Alice's belief tends towards the wrong type temporarily, she starts increasingly revealing both attributes. As her belief converges to 100%. Which of the filters is optimal depends on the state \mathbf{X}_t , of course, so Alice uses all three filters under the optimal policy.

6 Conclusion

In this paper, we have studied the information design problem of a sender who filters the information accessible to a boundedly rational receiver, in order to nudge or persuade them. We have introduced a rational inattention model of multi-attribute choice and provided an analytical solution (Theorem 1) whose conditional choice probabilities parsimoniously reflect the benefit and cost of acquiring information. We have formulated the dynamic problem of the sender as a partially observable Markov decision process (POMDP) and solved it numerically using partially observable Monte Carlo planning (POMCP). Simulations demonstrate that the sender succeeds in learning an information policy that is effective in inducing desirable receiver choices.

Our findings have implications beyond economics and marketing, for example for the design of recommender systems: The sender's choice of an action selection matrix \mathbf{A}_t can be viewed as a recommender system determining the subset of available options (and the order in which they are presented). The feature selection of \mathbf{F}_t amounts to information design.

There are numerous avenues for future research. First, it could be insightful to extend

our model to allow for a strategically sophisticated receiver taking into account the sender's motives in choosing an information filter. Second, considering alternative information cost functions for the receiver would be important in several respects: This could include cost functions that rationalize heuristic strategies attending to a subset of attributes only, for example, but also cost functions that are sensitive to the order in which options and attributes are presented. Lastly, relaxing the assumption of free disposal of information would open many more possibilities for the sender to help the receiver make better choices.

References

- David J Aldous, Illdar A Ibragimov, and Jean Jacod. École d'Été de Probabilités de Saint-Flour XIII — 1983. Lecture Notes in Mathematics, pages 1–198, 2006. ISSN 0075-8434. doi: 10.1007/bfb0099421.
- Guy Aridor, Duarte Goncalves, Daniel Kluver, Ruoyan Kong, and Joseph Konstan. The Economics of Recommender Systems: Evidence from a Field Experiment on MovieLens. arXiv, 2022. doi: 10.48550/arxiv.2211.14219.
- Dirk Bergemann and Stephen Morris. Information Design: A Unified Perspective. Journal of Economic Literature, 57(1):44–95, 2019. ISSN 0022-0515. doi: 10.1257/jel.20181489.
- Ann Blair. Information Overload's 2,300-Year-Old History. Harvard Business Review, 2011.
- Alexander W. Bloedel and Ilya Segal. Persuading a Rationally Inattentive Agent. *Working Paper*, 2021.
- Stefan Bucher and Andrew Caplin. Inattention and Inequity in School Matching. NBER WP, 2021.
- Jerome R. Busemeyer, Sebastian Gluth, Jörg Rieskamp, and Brandon M. Turner. Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-based Decisions. *Trends in Cognitive Sciences*, 23(3):251–263, 2019. ISSN 1364-6613. doi: 10.1016/j.tics.2018.12.003.
- Alison Buttenheim, Robert Moffitt, and Alexandra Beatty, editors. Behavioral Economics: Policy Impact and Future Directions. The National Academies Press, Washington, DC, 2023. doi: 10.17226/26874.
- Frederick Callaway, Mathew Hardy, and Thomas L. Griffiths. Optimal nudging for cognitively bounded agents: A framework for modeling, predicting, and controlling the effects of choice architectures. *PsyArXiv*, 2022.

- Andrew Caplin. The Science of Mistakes, Lecture Notes on Economic Data Engineering. World Scientific Lecture Notes in Economics and Policy, 2023. ISSN 2630-4872. doi: 10.1142/9789811262395\ 0005.
- Andrew Caplin and Daniel Martin. Comparison of Decisions under Unknown Experiments. Journal of Political Economy, 129(11):3185–3205, 2021. ISSN 0022-3808. doi: 10.1086/ 716104.
- Andrew Caplin, Mark Dean, and John Leahy. Rational Inattention, Optimal Consideration Sets, and Stochastic Choice. *Review of Economic Studies*, 86:1061—1094, 2019. doi: 10.1093/restud/rdy037/5060717.
- Andrew Caplin, Daniel Martin, and Philip Marx. Revealed Bayesian Learning. *Working Paper*, 2022.
- Alexander Chernev, Ulf Böckenholt, and Joseph Goodman. Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, 25(2):333–358, 2015. ISSN 1057-7408. doi: 10.1016/j.jcps.2014.08.002.
- Michael Choi, Anovia Yifan Dai, and Kyungmin Kim. Consumer Search and Price Competition. *Econometrica*, 86(4):1257–1281, 2018. ISSN 0012-9682. doi: 10.3982/ecta14837.
- Giovanni Compiani, Gregory Lewis, Sida Peng, and Will Wang. Online Search and Product Rankings: A Double Index Approach. SSRN Electronic Journal, 2021. doi: 10.2139/ssrn. 3898134.
- Mark Dean, Dilip Ravindran, and Jörg Stoye. A Better Test of Choice Overload. *arXiv*, 2022. doi: 10.48550/arxiv.2212.03931.
- Stefano DellaVigna and Elizabeth Linos. RCTs to Scale: Comprehensive Evidence From Two Nudge Units. *Econometrica*, 90(1):81–116, 2022. ISSN 0012-9682. doi: 10.3982/ecta18709.
- Ambuj Dewan. Costly Multidimensional Information. Working Paper, 2020.
- Shaddin Dughmi and Haifeng Xu. Algorithmic Bayesian Persuasion. SIAM Journal on Computing, 50(3):STOC16–68–STOC16–97, 2019. ISSN 0097-5397. doi: 10.1137/16m1098334.
- Maxim Egorov, Zachary N. Sunberg, Edward Balaban, Tim A. Wheeler, Jayesh K. Gupta, and Mykel J. Kochenderfer. POMDPs.jl: A Framework for Sequential Decision Making under Uncertainty. *Journal of Machine Learning Research*, 18(26):1–5, 2017.

- Benjamin Enke and Thomas Graeber. Cognitive Uncertainty. NBER Working Paper, 26518, 2021.
- Daniel Fleder and Kartik Hosanagar. Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science*, 55(5):697–712, 2009. ISSN 0025-1909. doi: 10.1287/mnsc.1080.0974.
- Cary Frydman and Lawrence J Jin. Efficient Coding and Risky Choice. *The Quarterly Journal of Economics*, 137(1):161–213, 2022. ISSN 0033-5533. doi: 10.1093/qje/qjab031.
- Paul Glimcher. Efficiently irrational: deciphering the riddle of human choice. Trends in Cognitive Sciences, 26(8):669–687, 2022.
- John R Hauser, Glen L Urban, Guilherme Liberali, and Michael Braun. Website Morphing. Marketing Science, 28(2):202–223, 2009a. ISSN 0732-2399. doi: 10.1287/mksc.1080.0459.
- John R Hauser, Glen L Urban, Guilherme Liberali, and Michael Braun. Rejoinder Response to Comments on "Website Morphing". Marketing Science, 28(2):227–228, 2009b. ISSN 0732-2399. doi: 10.1287/mksc.1080.0485.
- John R Hauser, Guilherme (Gui) Liberali, and Glen L Urban. Website Morphing 2.0: Switching Costs, Partial Exposure, Random Exit, and When to Morph. Management Science, 60(6):1594–1616, 2014. ISSN 0025-1909. doi: 10.1287/mnsc.2014.1961.
- Harold V. Henderson. C85. On the inverse of a completely symmetric matrix. Journal of Statistical Computation and Simulation, 12(2):138–139, 1981. ISSN 0094-9655. doi: 10.1080/00949658108810443.
- Elisabeth Honka and Pradeep Chintagunta. Simultaneous or Sequential? Search Strategies in the U.S. Auto Insurance Industry. *Marketing Science*, 36(1):21–42, 2017. ISSN 0732-2399. doi: 10.1287/mksc.2016.0995.
- Elisabeth Honka, Ali Hortaçsu, and Matthijs Wildenbeest. Chapter 4 Empirical search and consideration sets. *Handbook of the Economics of Marketing*, 1:193–257, 2019. ISSN 2452-2619. doi: 10.1016/bs.hem.2019.05.002.
- Benjamin Hébert and Michael Woodford. Rational Inattention when Decisions Take Time. NBER Working Paper 26415, 2021. URL http://www.nber.org/papers/w26415.
- Benjamin Hébert and Weijie Zhong. Engagement Maximization. arXiv, 2022. doi: 10.48550/arxiv.2207.00685.

- Sheena S. Iyengar and Mark R. Lepper. When Choice is Demotivating: Can One Desire Too Much of a Good Thing? Journal of Personality and Social Psychology, 79(6):995–1006, 2000. ISSN 0022-3514. doi: 10.1037/0022-3514.79.6.995.
- Jacob Jacoby. Perspectives on Information Overload. *Journal of Consumer Research*, 10(4): 432, 1984. ISSN 0093-5301. doi: 10.1086/208981.
- Jacob Jacoby, Donald E. Speller, and Carol A. Kohn Berning. Brand Choice Behavior as a Function of Information Load. *Journal of Marketing Research*, 1(1):33–42, 1974.
- M James. The generalised inverse. *The Mathematical Gazette*, 62(420):109–114, 1978. ISSN 0025-5572. doi: 10.2307/3617665.
- Eric J. Johnson. The elements of choice: Why the way we decide matters. Riverhead Books, New York, NY, 2021. ISBN 9780593084434.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998. ISSN 0004-3702. doi: 10.1016/s0004-3702(98)00023-x.
- Emir Kamenica. Bayesian Persuasion and Information Design. Annual Review of Economics, 11(1):1–24, 2019. ISSN 1941-1383. doi: 10.1146/annurev-economics-080218-025739.
- Emir Kamenica and Matthew Gentzkow. Bayesian Persuasion. American Economic Review, 101(6):2590–2615, 2011. ISSN 0002-8282. doi: 10.1257/aer.101.6.2590.
- Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization. arXiv, 2022. doi: 10.48550/arxiv.2202.11776.
- Gui Liberali and Alina Ferecatu. Morphing for Consumer Dynamics: Bandits Meet Hidden Markov Models. Marketing Science, 41(4):341–366, 2022. ISSN 0732-2399. doi: 10.1287/ mksc.2021.1346.
- Xiao Liu. Dynamic Coupon Targeting Using Batch Deep Reinforcement Learning: An Application to Livestream Shopping. *Marketing Science*, 2022. ISSN 0732-2399. doi: 10.1287/mksc.2022.1403.
- Xiao Liu. Deep Learning in Marketing: A Review and Research Agenda. In K. Sudhir and Olivier Toubia, editors, Artificial Intelligence in Marketing, volume 20 of Review of Marketing Research, pages 239–271. Emerald Publishing Limited, Bingley, 2023. ISBN 978-1-80262-876-0.

- Graham Loomes and Robert Sugden. Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal*, 92(368):805, 1982. ISSN 0013-0133. doi: 10.2307/2232669.
- Robert E. Lucas. Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy, 1:19–46, 1976. ISSN 0167-2231. doi: 10.1016/s0167-2231(76) 80003-6.
- Maximilian Maier, František Bartoš, T D Stanley, David R Shanks, Adam J L Harris, and Eric-Jan Wagenmakers. No evidence for nudging after adjusting for publication bias. *Proceedings of the National Academy of Sciences*, 119(31):e2200300119, 2022. ISSN 0027-8424. doi: 10.1073/pnas.2200300119.
- Filip Matejka and Alisdair McKay. Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. American Economic Review, 105(1):272—298, 2015.
- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational Inattention: A Review. Journal of Economic Literature, 61(1):226–273, 2023. ISSN 0022-0515. doi: 10.1257/jel. 20211524.
- Daniel McFadden. Choice What Can Go Wrong? NBER Working Paper, No. 31165, 2023.
- Bryce McLaughlin and Jann Spiess. Algorithmic Assistance with Recommendation-Dependent Preferences. arXiv, 2022. doi: 10.48550/arxiv.2208.07626.
- Stephanie Mertens, Mario Herberz, Ulf J. J. Hahnel, and Tobias Brosch. The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1):e2107346118, 2022. ISSN 0027-8424. doi: 10.1073/pnas.2107346118.
- Jianjun Miao, Jieran Wu, and Eric R. Young. Multivariate Rational Inattention. Econometrica, 90(2):907–945, 2022. ISSN 0012-9682. doi: 10.3982/ecta18086.
- Adam Sanjurjo. Search with multiple attributes: Theory and empirics. Games and Economic Behavior, 104:535–562, 2017. ISSN 0899-8256. doi: 10.1016/j.geb.2017.05.009.
- Babur De los Santos, Ali Hortaçsu, and Matthijs R Wildenbeest. Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior. American Economic Review, 102(6):2955–2980, 2012. ISSN 0002-8282. doi: 10.1257/aer.102.6.2955.

- Benjamin Scheibehenne, Rainer Greifeneder, and Peter M Todd. Can There Ever Be Too Many Options? A Meta-Analytic Review of Choice Overload. *Journal of Consumer Research*, 37(3):409–425, 2010. ISSN 0093-5301. doi: 10.1086/651235.
- Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020. ISSN 0028-0836. doi: 10.1038/s41586-020-03051-4.
- David Silver and Joel Veness. Monte-Carlo planning in large POMDPs. Advances in Neural Information Processing Systems, 23, 2010.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354– 359, 2017. ISSN 0028-0836. doi: 10.1038/nature24270.
- Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003. ISSN 0304-3932. doi: 10.1016/s0304-3932(03)00029-1.
- George J. Stigler. The Economics of Information. *Journal of Political Economy*, 69(3): 213-225, 1961.
- Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions about Health, Wealth, and Happiness.* Yale University Press, New Haven, 2008.
- Raluca Ursu, Stephan Seiler, and Elisabeth Honka. The Sequential Search Model: A Framework for Empirical Research. SSRN Electronic Journal, 2023. doi: 10.2139/ssrn.4357557.
- Martin L Weitzman. Optimal Search for the Best Alternative. *Econometrica*, 47(3):641, 1979. ISSN 0012-9682. doi: 10.2307/1910412.
- Michael Woodford. Modeling Imprecision in Perception, Valuation, and Choice. Annual Review of Economics, 12(1):579–601, 2020. ISSN 1941-1383. doi: 10.1146/ annurev-economics-102819-040518.
- Xiaozhi Yang and Ian Krajbich. A Dynamic Computational Model of Gaze and Choice in Multi-Attribute Decisions. *Psychological Review*, 130(1):52–70, 2023. ISSN 0033-295X. doi: 10.1037/rev0000350.

A Notation

Notation	Description	Definition
a	receiver action/option	$a \in \mathcal{A}$
${\mathcal A}$	grand set of receiver actions	
$A(\mathbf{A})$	receiver's choice set	$A(\mathbf{A}) := \{ a \in \mathcal{A} : \sum_{i} \mathbf{A}_{ia} > 0 \} \subseteq \mathcal{A}$
\mathbf{A}	action selection matrix	$\mathbf{A} \in \boldsymbol{\mathcal{A}} \ (ext{non-zero rows of } \hat{\mathbf{A}})$
$\hat{\mathbf{A}}$	full action selection matrix	$\hat{\mathbf{A}} \in \{0,1\}^{m imes m}: 1_m^T \hat{\mathbf{A}} \leq 1_m, \hat{\mathbf{A}} 1_m \leq 1_m^T$
${\cal A}$	set of action selection matrices	$\{\mathbf{A} \in \{0,1\}^{k imes m} : 1 \le k \le m, 1_k^T \mathbf{A} \le 1_m, \mathbf{A} 1_m =$
b	sender belief	$b \in \Delta(\Theta)$
$oldsymbol{e}_a$	standard basis vector	$oldsymbol{e} \in \{0,1\}^k,oldsymbol{e}_a=1,\sum_ioldsymbol{e}_i=1$
f	attribute/feature	$f\in \mathcal{F}$
${\cal F}$	set of attributes/features	
\mathbf{F}	feature selection matrix	$\mathbf{F} \in \boldsymbol{\mathcal{F}} \ (\text{non-zero rows of } \hat{\mathbf{F}})$
$\hat{\mathbf{F}}$	full feature selection matrix	$\hat{\mathbf{F}} \in \{0,1\}^{n imes n} : 1_n^T \hat{\mathbf{F}} \le 1_n, \hat{\mathbf{F}} 1_n \le 1_n^T$
${\cal F}$	set of feature selection matrices	$\{\mathbf{F} \in \{0,1\}^{l \times n} : 1 \le l \le n, 1_l^T \mathbf{F} \le 1_n, \mathbf{F} 1_n = 1_l^T\}$
(\mathbf{A},\mathbf{F})	information filter	$(\mathbf{A},\mathbf{F})\in\mathcal{H}$
${\cal H}$	set of information filters	${\mathcal H}:={\mathcal A} imes {\mathcal F}$
k	number of options in slate	$k = rank(\mathbf{A})$
$K_{\theta}(P; \mathbf{A}, \mathbf{F}, \mu)$	information cost function	$K: \mathcal{P} \times \Theta \times \mathcal{H} \times \Delta(\mathcal{X}) \to \bar{\mathbb{R}}$
l	number of attributes shown	$l = rank(\mathbf{F})$
m	number of actions	$m = \mathcal{A} $
n	number of attributes/features	$n = \mathcal{F} $
$O(a \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu)$	sender observations	$O: \mathcal{X} \times \Theta \times \mathcal{H} \times \Delta(\mathcal{X}) \to \Delta(A(\mathbf{A}))$
$P(a \mathbf{Y}; \theta, \mathbf{A}, \mathbf{F}, \mu)$	receiver choice model	$P: \mathcal{Y} \times \Theta \times \mathcal{H} \times \Delta(\mathcal{X}) \to \Delta(\mathcal{A})$
$\mathcal{P}(\mathbf{A},\mu)$	receiver strategy set	$\{P: P(a \mathbf{A}\mathbf{X}\mathbf{F}^T; \theta, \mathbf{A}, \mathbf{F}, \mu) = 0 \forall a \notin A(\mathbf{A})\}$
$Q(\mathbf{X}, b, \mathbf{A}, \mathbf{F}, \mu)$	sender's state-action value	
$R^v_{lpha}(a,\mathbf{X}, heta,\mathbf{A},\mathbf{F};\mu)$	sender's reward function, eq. 4	
S	sender's state space	$\mathcal{S} = \mathcal{X} imes \Theta$
t	time period	$t \in \{0, \dots, \infty\}$
$\mathcal{T}(\mathbf{X}', heta' \mathbf{X}, heta, \mathbf{A}, \mathbf{F})$	state transition function	$\mathcal{T}: \mathcal{X} \times \Theta imes \mathcal{H} o \Delta(\mathcal{X} imes \Theta)$
$u_{ heta}(a, \mathbf{Z})$	receiver utility	$u: \mathcal{A} \times \zeta^{k \times n} \times \Theta \to \mathbb{R}$
$U^v_{lpha}({f X}, heta,{f A},{f F};\mu)$	conditional welfare, eq. 5	$U^v_\alpha: \mathcal{X} \times \Theta \times \mathcal{H} \times \Delta(\mathcal{X}) \to \mathbb{R}$
$v(a, \mathbf{Z}, \theta)$	sender's (physical) utility	$v: \mathcal{A} \times \zeta^{k \times n} \times \Theta \to \mathbb{R}$
$V^{\pi}(\mathbf{X},b,\mu)$	value function of sender problem	$V^{\pi}: \mathcal{X} \times \Delta(\Theta) \times \Delta(\mathcal{X}) \to \mathbb{R}$
$oldsymbol{w}$	receiver's preference weights	$oldsymbol{w} \in \mathbb{R}^n$
X	state/payoff matrix	$\mathbf{X} \in \mathcal{X} = \zeta^{\mathcal{A} imes \mathcal{F}}$
$x_{a,f}$	value of attribute f of action a	$x_{a,f} \in \zeta$

		(T) m
$oldsymbol{x}$	state vector (row-major repres. of \mathbf{X})	$oldsymbol{x} = vec(\mathbf{X}^{T}) \in \zeta^{m \cdot n}$
Y	filtered state matrix	$\mathbf{Y} = \mathbf{A}\mathbf{X}\mathbf{F}^T \in \mathcal{Y} = \zeta^{k imes l}$
$oldsymbol{y}$	row-major repres. of filtered state	$oldsymbol{y} = vec(\mathbf{Y}^T) = (\mathbf{A} \otimes \mathbf{F}) oldsymbol{x} \in \zeta^{k \cdot l}$
lpha	sender's discount factor on info. costs	$\alpha \in [0,1]$
$oldsymbol{eta}$	receiver's effective preference weights	$oldsymbol{eta} := oldsymbol{w} / \kappa$
γ	sender's temporal discount factor	$\gamma \in [0,1)$
ζ	set of elementary prizes	
$ heta = (oldsymbol{w},\kappa)$	receiver type	$ heta\in\Theta$
κ	marginal cost per bit of information	$\kappa \in \mathbb{R}_0^+$
$\mu(\mathbf{X})$	receiver's prior belief	$\mu \in \Delta(\mathcal{X})$
$\pi(\mathbf{A},\mathbf{F} \mathbf{X},b,\mu)$	sender policy	$\pi: \mathcal{X} \times \Delta(\Theta) \times \Delta(\mathcal{X}) \to \Delta(\mathcal{H})$
au(heta' heta)	receiver type distribution	$\tau: \Theta \to \Delta(\Theta)$

B Proofs

B.1 Proof of Lemma 1

Let $\mathbf{x} \equiv vec(\mathbf{X}^T)$ be the (row-major) vectorization of matrix \mathbf{X} .⁶ It follows from $vec(\mathbf{ABC}) = (\mathbf{C}^T \otimes \mathbf{A})vec(\mathbf{B})$ that

$$\boldsymbol{y} := vec(\mathbf{Y}^T) = vec(\mathbf{F}\mathbf{X}^T\mathbf{A}^T) = (\mathbf{A}\otimes\mathbf{F})\mathbf{X}^T =: \mathbf{H}\boldsymbol{x}$$

where $\mathbf{H} := \mathbf{A} \otimes \mathbf{F}$. Let $\mathbf{H}^+ = \mathbf{H}^T (\mathbf{H}\mathbf{H}^T)^{-1}$ be the (right) Moore-Penrose inverse of (full row-rank) matrix \mathbf{H} . It is a standard result in linear algebra (e.g. James, 1978) that if the linear system of equations $\mathbf{y} = \mathbf{H}\mathbf{x}$ has any solutions (which is the case if and only if $\mathbf{H}^+\mathbf{H}\mathbf{y} = \mathbf{y}$), then they are all given by the set of vectors \mathbf{x}

$$\{\mathbf{H}^+ oldsymbol{y} + (\mathbf{I}_{m \cdot n} - \mathbf{H}^+ \mathbf{H}) oldsymbol{z} : oldsymbol{z} \in \mathbb{R}^{m \cdot n}\}$$

spanned by any $\boldsymbol{z} \in \mathbb{R}^{m \cdot n}$. The result follows immediately.

B.2 Pseudoinverse of A

Note that any $\mathbf{A} \in \mathcal{A}$ has full row-rank and is semi-orthogonal $(\mathbf{A}\mathbf{A}^T = \mathbf{I})$ so that

$$\mathbf{A}^+ := \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} = \mathbf{A}^T$$

⁶We adopt the common convention that $vec(\mathbf{X})$ denotes the column-major vectorization, so that $vec(\mathbf{X}^T)$ is the row-major vectorization of \mathbf{X} .

is its right (Moore-Penrose) pseudoinverse: $\mathbf{A}\mathbf{A}^+ = \mathbf{I}_k$.

Lemma 5. Given any $k \times k$ permutation matrix \mathbf{P}_k , let $\mathbf{P}_m^{\mathbf{P}_k} := \mathbf{A}^+ \mathbf{P}_k \mathbf{A} + \mathbf{I}_m - \mathbf{A}^+ \mathbf{A}$. Then $\mathbf{P}_m^{\mathbf{P}_k}$ is a $m \times m$ permutation matrix.⁷ Further, it is the case that

$$\mathbf{A}\mathbf{P}_{m}^{\mathbf{P}_{k}} = \mathbf{P}_{k}\mathbf{A} \tag{18}$$

and thus $\mathbf{A}(\mathbf{P}_m^{\mathbf{P}_k}\mathbf{X})\mathbf{F}^T = \mathbf{P}_k(\mathbf{A}\mathbf{X}\mathbf{F}^T)$.

Proof. Note that

$$\mathbf{A}\mathbf{P}_m^{\mathbf{P}_k} = \mathbf{A}\mathbf{A}^+\mathbf{P}_k\mathbf{A} + \mathbf{A}\mathbf{I}_m - \mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{I}_k\mathbf{P}_k\mathbf{A} + \mathbf{A}\mathbf{I}_m - \mathbf{I}_k\mathbf{A} = \mathbf{P}_k\mathbf{A}.$$

B.3 Proof of Lemma 2

For any $k \times k$ permutation matrix \mathbf{P}_k and all $\mathbf{Y} \in \mathcal{Y}$, it is the case that

$$\mu_{\mathbf{Y}}(\mathbf{P}_{k}\mathbf{Y}|\mathbf{A},\mathbf{F}) = \int_{\{\mathbf{Z}:\mathbf{A}\mathbf{Z}\mathbf{F}^{T}=\mathbf{P}_{k}\mathbf{Y}\}} d\mu(\mathbf{Z})$$

$$= \int_{\{\mathbf{X}:\mathbf{A}\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X}\mathbf{F}^{T}=\mathbf{P}_{k}\mathbf{Y}\}} d\mu(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X})$$

$$= \int_{\{\mathbf{X}:\mathbf{P}_{k}\mathbf{A}\mathbf{X}\mathbf{F}^{T}=\mathbf{P}_{k}\mathbf{Y}\}} d\mu(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X})$$

$$= \int_{\{\mathbf{X}:\mathbf{A}\mathbf{X}\mathbf{F}^{T}=\mathbf{Y}\}} d\mu(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X})$$

$$= \int_{\{\mathbf{X}:\mathbf{A}\mathbf{X}\mathbf{F}^{T}=\mathbf{Y}\}} d\mu(\mathbf{X})$$

$$= \mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A},\mathbf{F})$$
(8)

where the first equality follows by the definition of equation 7, the second equality from substituting $\mathbf{Z} = \mathbf{P}_m^{\mathbf{P}_k} \mathbf{X}$, the third equality from equation 18, the fourth equality from \mathbf{P}_k being a permutation matrix and hence invertible, the fifth equality by Definition 1 from the fact that $\mathbf{P}_m^{\mathbf{P}_k}$ is a permutation matrix, and the last equality again from equation 7.

Moreover, for any $k \times k$ permutation matrix \mathbf{P}_k , all $\mathbf{Y} \in \mathcal{Y}$ and all $\mathbf{X} \in \chi(\mathbf{Y}; \mathbf{A}, \mathbf{F})$,

$$\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X}|\mathbf{P}_{k}\mathbf{Y},\mathbf{A},\mathbf{F}) = \frac{\mu_{\mathbf{X}}(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X})}{\mu_{\mathbf{Y}}(\mathbf{P}_{k}\mathbf{Y}|\mathbf{A},\mathbf{F})} = \frac{\mu_{\mathbf{X}}(\mathbf{P}_{m}^{\mathbf{P}_{k}}\mathbf{X})}{\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A},\mathbf{F})} = \frac{\mu_{\mathbf{X}}(\mathbf{X})}{\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A},\mathbf{F})} = \mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y},\mathbf{A},\mathbf{F})$$
(9)

where the first equality follows from equation 6, the second equality from equation 8, the third

⁷A permutation matrix is a square binary matrix with exactly one entry of 1 in each row and each column.

equality by Definition 1 from the fact that $\mathbf{P}_m^{\mathbf{P}_k}$ is an $m \times m$ permutation matrix, and the last equality again from equation 6.

B.4 Proof of Theorem 1

Suppressing all time subscripts in this proof, the expected utility conditional on \mathbf{Y} is given by

$$\tilde{u}_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) := \int_{\mathcal{X}} u_{\theta}(a, \mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F}) = \frac{\int_{\chi(\mathbf{Y}; \mathbf{A}, \mathbf{F})} u_{\theta}(a, \mathbf{A}\mathbf{X}) d\mu(\mathbf{X})}{\int_{\chi(\mathbf{Y}; \mathbf{A}, \mathbf{F})} d\mu(\mathbf{X})}$$
(19)

where the second equality follows from equation 6. The problem of the receiver (eq. 3) is thus equivalent to

$$\max_{P} \int_{\mathcal{Y}} \sum_{a \in A(\mathbf{A})} P(a|\mathbf{Y}; \theta, \mathbf{A}, \mathbf{F}, \mu) \tilde{u}_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) - \kappa I_{P}(a; vec(\mathbf{Y}))$$
(20)

because

$$\begin{split} &\int_{\mathcal{Y}} P(a|\mathbf{Y};\theta,\mathbf{A},\mathbf{F},\mu) \int_{\mathcal{X}} u_{\theta}(a,\mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y},\mathbf{A},\mathbf{F}) d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A},\mathbf{F}) \\ &= \int_{\mathcal{X}} P(a|\mathbf{A}\mathbf{X}\mathbf{F}^{T};\theta,\mathbf{A},\mathbf{F},\mu) u_{\theta}(a,\mathbf{A}\mathbf{X}) d\mu(\mathbf{X}). \end{split}$$

The transformed problem of equation 20 is a standard rational inattention problem in the state \mathbf{Y} , which we can solve by adapting the solution of Bucher and Caplin (2021) to the case of a matrix state with a row-exchangeable prior distribution.

To do so, fixing **A** and **F**, define a partition $supp(\mu_{\mathbf{Y}}) = \bigcup_{i \in I} \mathcal{Y}_i$ with $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}_i$ for some *i* if and only if there exist a permutation $\varrho : A(\mathbf{A}) \to A(\mathbf{A})$ and associated $k \times k$ permutation matrix \mathbf{P}_k^{ϱ} such that (i) $\mathbf{Y}' = \mathbf{P}_k^{\varrho} \mathbf{Y}$ and (ii) \mathbf{P}_k^{ϱ} is consistent with ϱ ($\boldsymbol{e}_{\varrho(a)} = \mathbf{P}_k^{\varrho} \boldsymbol{e}_a$).

We prove two auxiliary Lemmas before proceeding to the main proof.

B.4.1 Auxiliary Lemma 6

Lemma 6. Given any \mathbf{A}, \mathbf{F} , any row-exchangeable μ , and any $i \in I$,

$$\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) = \sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}'; \mathbf{A}, \mathbf{F}, \mu) \qquad \forall \mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}_i.$$

Proof. We have, for any $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}_i$ and any $c \in A(\mathbf{A})$, that

$$\widetilde{u}_{\theta}(\varrho(c), \mathbf{Y}'; \mathbf{A}, \mathbf{F}, \mu) = \int_{\mathcal{X}} u_{\theta}(\varrho(c), \mathbf{A}\mathbf{X}') d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}'|\mathbf{Y}', \mathbf{A}, \mathbf{F})
= \int_{\mathcal{X}} u_{\theta}(\varrho(c), \mathbf{A}\mathbf{X}') d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}'|\mathbf{P}_{k}^{\varrho}\mathbf{Y}, \mathbf{A}, \mathbf{F})
= \int_{\mathcal{X}} u_{\theta}(\varrho(c), \mathbf{A}\mathbf{P}_{m}^{\mathbf{P}_{k}^{\varrho}}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{P}_{m}^{\mathbf{P}_{k}^{\varrho}}\mathbf{X}|\mathbf{P}_{k}^{\varrho}\mathbf{Y}, \mathbf{A}, \mathbf{F})
= \int_{\mathcal{X}} u_{\theta}(\varrho(c), \mathbf{P}_{k}^{\varrho}\mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{P}_{m}^{\mathbf{P}_{k}^{\varrho}}\mathbf{X}|\mathbf{P}_{k}^{\varrho}\mathbf{Y}, \mathbf{A}, \mathbf{F})
= \int_{\mathcal{X}} u_{\theta}(c, \mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F})
= \widetilde{u}_{\theta}(c, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu)$$
(21)

where the first equality holds by equation 19, the second equality follows from the existence of a \mathbf{P}_{k}^{ϱ} such that $\mathbf{Y}' = \mathbf{P}_{k}^{\varrho}\mathbf{Y}$, the third equality from a change of variable (with the range remaining unchanged because $supp(\mu_{\mathbf{X}|\mathbf{Y}})$ is closed under permutation, by its row-exchangeability (eq. 9)), the fourth equality follows from equation 18, and the fifth equality from equations 2 and 9. The result then follows from

$$\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) = \sum_{c \in A(\mathbf{A})} \exp\left(\frac{1}{\kappa} \tilde{u}_{\theta}(c, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu)\right) = \sum_{c \in A(\mathbf{A})} \exp\left(\frac{1}{\kappa} \tilde{u}_{\theta}(\varrho(c), \mathbf{Y}'; \mathbf{A}, \mathbf{F}, \mu)\right)$$
$$= \sum_{c \in A(\mathbf{A})} \exp\left(\frac{1}{\kappa} \tilde{u}_{\theta}(c, \mathbf{Y}'; \mathbf{A}, \mathbf{F}, \mu)\right) = \sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}'; \mathbf{A}, \mathbf{F}, \mu)$$

where the second equality follows from eq. 21 and the third equality from summing over the entire domain of permutation ρ .

B.4.2 Auxiliary Lemma 7

Lemma 7. Given any \mathbf{A}, \mathbf{F} , any row-exchangeable μ , and any $i \in I$,

$$\int_{\mathcal{Y}_i} z_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) d\mathbf{Y} = \int_{\mathcal{Y}_i} z_{\theta}(b, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) d\mathbf{Y} \qquad \forall a, b \in A(\mathbf{A}).$$

Proof. Fix any permutation $\varrho : A(\mathbf{A}) \to A(\mathbf{A})$ with $b = \varrho(a)$ and define $\sigma_{\varrho} : \mathcal{Y} \to \mathcal{Y}$ for any $\mathbf{Y} \in \mathcal{Y}$ as $\sigma_{\varrho}(\mathbf{Y}) = \mathbf{P}_{k}^{\varrho}\mathbf{Y}$. We now show that σ_{ϱ} is a permutation on $supp(\mu_{\mathbf{Y}})$. To see this, first note that $\sigma_{\varrho}^{-1}(\mathbf{Y}') = \mathbf{P}_{k}^{\varrho T}\mathbf{Y}'$. Because the row-exchangeability of $\mu_{\mathbf{Y}}$ (eq. 8) implies that $supp(\mu_{\mathbf{Y}})$ is closed under permutation with any \mathbf{P}_{k} , it follows that $\sigma_{\varrho}^{-1}(\mathbf{Y}') \in supp(\mu_{\mathbf{Y}})$ for any $\mathbf{Y}' \in supp(\mu_{\mathbf{Y}})$, so σ_{ϱ} is surjective. Since $\sigma_{\varrho}^{-1}(\mathbf{Y}')$ is unique, σ_{ϱ} is also injective and hence bijective and a permutation. Because the definition of the partition $\{\mathcal{Y}_i\}$ implies that the image of σ_{ϱ} satisfies $\sigma_{\varrho}[\mathcal{Y}_i] \subseteq \mathcal{Y}_i$ for any $i \in I$, the restriction of σ_{ϱ} to \mathcal{Y}_i is in fact a permutation $\sigma_{\varrho}|_{\mathcal{Y}_i} : \mathcal{Y}_i \to \mathcal{Y}_i$ on \mathcal{Y}_i . It then follows that, for all $a, b \in A(\mathbf{A})$,

$$\begin{split} \int_{\mathcal{Y}_i} z_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) d\mathbf{Y} &= \int_{\mathcal{Y}_i} \exp\left(\frac{1}{\kappa} \tilde{u}_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu)\right) d\mathbf{Y} \\ &= \int_{\mathcal{Y}_i} \exp\left(\frac{1}{\kappa} \tilde{u}_{\theta}(\varrho(a), \sigma_{\varrho}|_{\mathcal{Y}_i}(\mathbf{Y}); \mathbf{A}, \mathbf{F}, \mu)\right) d\mathbf{Y} \\ &= \int_{\mathcal{Y}_i} \exp\left(\frac{1}{\kappa} \tilde{u}_{\theta}(\varrho(a), \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu)\right) d\mathbf{Y} \\ &= \int_{\mathcal{Y}_i} z_{\theta}(b, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu) d\mathbf{Y} \end{split}$$

where the second equality follows from equation 21, the third equality from summing over the whole domain of permutation $\sigma_{\varrho}|_{\mathcal{Y}_i}$, and the last equality from $b = \varrho(a)$.

B.4.3 Proof of Theorem 1

We now proceed to the proof of Theorem 1, suppressing the latter arguments of $z_{\theta}(a, \mathbf{Y}; \mathbf{A}, \mathbf{F}, \mu)$ for brevity. Lemma 6 implies that

$$Z_{\theta}(\mathbf{Y}) := \sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}) = \sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}') = Z_{\theta}(\mathbf{Y}') =: Z_{i} \qquad \forall \mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}_{i}$$
(22)

and equation 8 implies that $\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) = \mu_{\mathbf{Y}}(\mathbf{Y}'|\mathbf{A}, \mathbf{F}) =: \mu_i$ for all $\mathbf{Y}, \mathbf{Y}' \in \mathcal{Y}_i$. It follows, given some $i \in I$, that, for all $a, b \in A(\mathbf{A})$,

$$\int_{\mathcal{Y}_i} \mu_{\mathbf{Y}}(\mathbf{Y}) \frac{z_{\theta}(a, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mathbf{Y} = \frac{\mu_i}{Z_i} \int_{\mathcal{Y}_i} z_{\theta}(a, \mathbf{Y}) d\mathbf{Y} = \frac{\mu_i}{Z_i} \int_{\mathcal{Y}_i} z_{\theta}(b, \mathbf{Y}) d\mathbf{Y} = \int_{\mathcal{Y}_i} \mu_{\mathbf{Y}}(\mathbf{Y}) \frac{z_{\theta}(b, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mathbf{Y} \quad (23)$$

where the second equality follows from Lemma 7. This implies, for any $a, b \in A(\mathbf{A})$, that

$$\int_{\mathcal{Y}} \frac{z_{\theta}(a, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y}) = \int_{I} \int_{\mathcal{Y}_{i}} \frac{z_{\theta}(a, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y}) di = \int_{I} \int_{\mathcal{Y}_{i}} \frac{z_{\theta}(b, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y}) di = \int_{\mathcal{Y}} \frac{z_{\theta}(b, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} \mu_{\mathbf{Y}}(\mathbf{Y}) d\mathbf{Y} d\mu_{\mathbf{Y}}(\mathbf{Y}) d\mathbf{Y} d$$

so that

$$\int_{\mathcal{Y}} \frac{z_{\theta}(a, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y}) = 1/|A(\mathbf{A})| \qquad \forall a \in A(\mathbf{A}).$$
(25)

Any symmetric strategy P, i.e. that satisfies $\int_{\mathcal{Y}} P(c, \mathbf{Y}) d\mathbf{Y} = 1/|A(\mathbf{A})|$ for all $c \in A(\mathbf{A})$, therefore satisfies

$$\int_{\mathcal{Y}} \frac{z_{\theta}(a, \mathbf{Y}) \mu_{\mathbf{Y}}(\mathbf{Y})}{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y}) \int_{\mathcal{Y}} P(c, \mathbf{Y}) d\mathbf{Y}} d\mathbf{Y} = |A(\mathbf{A})| \int_{\mathcal{Y}} \frac{z_{\theta}(a, \mathbf{Y})}{Z_{\theta}(\mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y}) = 1 \qquad \forall a \in A(\mathbf{A})$$
(26)

where the first equality follows from the fact that for a symmetric strategy $\int_{\mathcal{Y}} P(c, \mathbf{Y}) d\mathbf{Y} = 1/|A(\mathbf{A})|$ for all $c \in A(\mathbf{A})$ by definition. The necessary and sufficient conditions for optimality (Caplin et al., 2019, Proposition 1) then imply that a symmetric strategy is a solution if and only if the statedependent choice probabilities satisfy the necessary conditions of Matejka and McKay (2015), which in this case reduce to equation 10.

B.5 Proof of Lemma 3

For any **A** and any $k \times k$ permutation matrix \mathbf{P}_k^{ϱ} with associated permutation $\varrho : A(\mathbf{A}) \to \mathbf{A}$, it is the case, for all $a \in A(\mathbf{A})$ and all $\mathbf{X} \in \mathcal{X}$, that

$$u_{\theta}(\varrho(a), \mathbf{P}_{k}^{\varrho} \mathbf{A} \mathbf{X}) = \boldsymbol{e}_{\varrho(a)}^{T} \mathbf{P}_{k}^{\varrho} \mathbf{A} \mathbf{X} \boldsymbol{w} = \boldsymbol{e}_{a}^{T} \mathbf{A} \mathbf{X} \boldsymbol{w} = u_{\theta}(a, \mathbf{A} \mathbf{X})$$
(2)

where the first equality follows by the definition of linear utility (eq. 11) and the second equality from $\boldsymbol{e}_{\varrho(a)}^T \mathbf{P}_k^{\varrho} = ((\mathbf{P}_k^{\varrho})^T \boldsymbol{e}_{\varrho(a)})^T = ((\mathbf{P}_k^{\varrho})^{-1} \boldsymbol{e}_{\varrho(a)})^T = \boldsymbol{e}_a^T$.

B.6 Proof of Lemma 4

To see that $\mathbf{M} = \mathbf{1}_m \otimes \boldsymbol{m}^T \in \mathbb{R}^{m \times n}$ and $\mathbf{U} = c \mathbf{1}_m \mathbf{1}_m^T + (\nu - c) \mathbf{I}_m \in \mathbb{R}^{m \times m}$ are sufficient for a matrix Gaussian distribution to be row-exchangeable, note that

$$\mu(\mathbf{P}_{m}\mathbf{X}|\mathbf{M},\mathbf{U},\mathbf{V}) = \frac{\exp\left(-\frac{1}{2}tr[\mathbf{V}^{-1}(\mathbf{P}_{m}\mathbf{X}-\mathbf{M})^{T}\mathbf{U}^{-1}(\mathbf{P}_{m}\mathbf{X}-\mathbf{M})]\right)}{(2\pi)^{mn/2}|\mathbf{V}|^{m/2}|\mathbf{U}|^{n/2}}$$
$$= \frac{\exp\left(-\frac{1}{2}tr[\mathbf{V}^{-1}(\mathbf{X}-\mathbf{M})^{T}\mathbf{P}_{m}^{T}\mathbf{U}^{-1}\mathbf{P}_{m}(\mathbf{X}-\mathbf{M})]\right)}{(2\pi)^{mn/2}|\mathbf{V}|^{m/2}|\mathbf{U}|^{n/2}}$$
$$= \frac{\exp\left(-\frac{1}{2}tr[\mathbf{V}^{-1}(\mathbf{X}-\mathbf{M})^{T}\mathbf{U}^{-1}(\mathbf{X}-\mathbf{M})]\right)}{(2\pi)^{mn/2}|\mathbf{V}|^{m/2}|\mathbf{U}|^{n/2}}$$
$$= \mu(\mathbf{X}|\mathbf{M},\mathbf{U},\mathbf{V})$$

where the first equality is the probability density function of a matrix Gaussian distribution, the second equality follows from \mathbf{M} consisting of identical rows so that $\mathbf{P}_m \mathbf{M} = \mathbf{M}$, and the third equality from $\mathbf{P}_m^T \mathbf{U}^{-1} \mathbf{P}_m = \mathbf{U}^{-1}$ which holds because $\mathbf{U}^{-1} = \frac{1}{\nu-c} (\mathbf{I}_m - \frac{c}{\nu+(m-1)c} \mathbf{1}_m \mathbf{1}_m^T)$ (Henderson, 1981), which is a completely symmetric matrix.

B.7 Auxiliary Lemma

Lemma 8. Let $\overline{\mathbf{Y}} = \mathbf{A}\mathbf{X}\overline{\mathbf{F}}^T$ where $\overline{\mathbf{F}}$ consists of the non-zero rows of matrix $\mathbf{I}_n - diag(\mathbf{1}_n^T \mathbf{F})$ (if any; otherwise $\overline{\mathbf{F}} = \mathbf{0}$), collecting all features of $\mathbf{A}\mathbf{X}$ not present in \mathbf{Y} . The Gaussian prior of equation 12 implies the conditional distribution

$$vec(\overline{\mathbf{Y}})|vec(\mathbf{Y}) \sim \mathcal{N}(\boldsymbol{m}_{\overline{\boldsymbol{y}}|\boldsymbol{y}}, \boldsymbol{\Sigma}_{\overline{\boldsymbol{y}}|\boldsymbol{y}})$$

with conditional mean

$$m_{\overline{y}|y} = vec \left(\mathbf{A} \mathbf{M} \overline{\mathbf{F}}^T + (\mathbf{Y} - \mathbf{A} \mathbf{M} \mathbf{F}^T) (\mathbf{F} \mathbf{V} \mathbf{F}^T)^{-1} \mathbf{F} \mathbf{V} \overline{\mathbf{F}}^T
ight).$$

Proof. Letting $\tilde{\mathbf{F}}^T = \begin{pmatrix} \mathbf{F}^T & \overline{\mathbf{F}}^T \end{pmatrix}$, we have

$$\mathbf{A}\mathbf{X}\tilde{\mathbf{F}}^T \sim \mathcal{M}\mathcal{N}_{m \times n} \left(\mathbf{A}\mathbf{M}\tilde{\mathbf{F}}^T, \mathbf{A}\mathbf{U}\mathbf{A}^T, \tilde{\mathbf{F}}\mathbf{V}\tilde{\mathbf{F}}^T\right)$$

which implies that 8

$$vec\left(\mathbf{A}\mathbf{X}\tilde{\mathbf{F}}^{T}
ight) = \begin{pmatrix} (\mathbf{F}\otimes\mathbf{A})vec(\mathbf{X})\\ (\overline{\mathbf{F}}\otimes\mathbf{A})vec(\mathbf{X}) \end{pmatrix} = \begin{pmatrix} vec(\mathbf{A}\mathbf{X}\mathbf{F}^{T})\\ vec(\mathbf{A}\mathbf{X}\overline{\mathbf{F}}^{T}) \end{pmatrix} \sim \mathcal{N}\left(\boldsymbol{m}_{\tilde{\boldsymbol{y}}}, \boldsymbol{\Sigma}_{\tilde{\mathbf{Y}}}
ight)$$

with

$$oldsymbol{m}_{ ilde{oldsymbol{y}}} = vec\left(\mathbf{A}\mathbf{M} ilde{\mathbf{F}}^T
ight) = egin{pmatrix} vec(\mathbf{A}\mathbf{M}\mathbf{F}^T) \ vec(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^T) \end{pmatrix}$$

and

$$\boldsymbol{\Sigma}_{\tilde{\mathbf{Y}}} = \left(\tilde{\mathbf{F}}\mathbf{V}\tilde{\mathbf{F}}^{T}\right) \otimes \left(\mathbf{A}\mathbf{U}\mathbf{A}^{T}\right) = \begin{pmatrix} (\mathbf{F}\mathbf{V}\mathbf{F}^{T}) \otimes (\mathbf{A}\mathbf{U}\mathbf{A}^{T}) & (\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^{T}) \otimes (\mathbf{A}\mathbf{U}\mathbf{A}^{T}) \\ (\overline{\mathbf{F}}\mathbf{V}\mathbf{F}^{T}) \otimes (\mathbf{A}\mathbf{U}\mathbf{A}^{T}) & (\overline{\mathbf{F}}\mathbf{V}\overline{\mathbf{F}}^{T}) \otimes (\mathbf{A}\mathbf{U}\mathbf{A}^{T}) \end{pmatrix}.$$

It follows that⁹

$$\begin{split} \boldsymbol{m}_{\overline{\boldsymbol{y}}|\boldsymbol{y}} &\equiv \mathbb{E}[vec(\mathbf{A}\mathbf{X}\overline{\mathbf{F}}^{T})|vec(\mathbf{Y})] \\ &= vec(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T}) + \left((\overline{\mathbf{F}}\mathbf{V}\mathbf{F}^{T})\otimes(\mathbf{A}\mathbf{U}\mathbf{A}^{T})\right)\left((\mathbf{F}\mathbf{V}\mathbf{F}^{T})\otimes(\mathbf{A}\mathbf{U}\mathbf{A}^{T})\right)^{-1}(vec(\mathbf{Y}) - vec(\mathbf{A}\mathbf{M}\mathbf{F}^{T})) \\ &= vec(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T}) + \left((\overline{\mathbf{F}}\mathbf{V}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\right)\otimes\left((\mathbf{A}\mathbf{U}\mathbf{A}^{T})(\mathbf{A}\mathbf{U}\mathbf{A}^{T})^{-1}\right)(vec(\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})) \\ &= vec(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T}) + \left((\overline{\mathbf{F}}\mathbf{V}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\otimes\mathbf{I}_{m}\right)(vec(\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})) \\ &= vec(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T}) + vec((\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})((\overline{\mathbf{F}}\mathbf{V}\mathbf{F}^{T})^{-1})^{T}) \\ &= vec\left(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T} + (\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^{T}\right) \end{split}$$

B.8 Proof of Corollary 1

With the linear preferences of equation 11,

$$\int_{\mathcal{X}} u_{\theta}(a, \mathbf{A}\mathbf{X}) d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F}) = \int_{\mathcal{X}} e_a^T \mathbf{A}\mathbf{X} w d\mu_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F}) = e_a^T \mathbf{A}\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}[\mathbf{X}|\mathbf{Y}, \mathbf{A}, \mathbf{F}] w.$$

⁸Note that we are using column-major vectorization here, unlike in the rest of the paper, because we want features to be adjacent in the resulting vector.

⁹Note that if \mathbf{F} is square (i.e. $l = rank(\mathbf{F}) = n$), then $(\mathbf{F}\mathbf{V}\mathbf{F}^T)^{-1}(\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^T) = (\overline{\mathbf{F}}\mathbf{F}^{-1})^T$ and the conditional expectation reduces to $vec(\mathbf{Y}(\mathbf{F}^{-1})^T)$, reflecting the fact that if all features are present then \mathbf{Y} contains all information on \mathbf{X} .

The result then follows immediately from Theorem 1 along with the fact that the expected utility conditional on \mathbf{Y} is

$$\begin{split} \tilde{u}_{\theta}(a,\mathbf{Y};\mathbf{A},\mathbf{F},\mu) &= \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}[u_{\theta}(a,\mathbf{A}\mathbf{X})|\mathbf{Y}] \\ &= \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}[vec(e_{a}^{T}\mathbf{A}\mathbf{X}w)|vec(\mathbf{Y})] \\ &= \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}[vec(e_{a}^{T}(\mathbf{A}\mathbf{X}\mathbf{F}^{T}\mathbf{F} + \mathbf{A}\mathbf{X}\overline{\mathbf{F}}^{T}\overline{\mathbf{F}})w)|vec(\mathbf{Y})] \\ &= (w^{T} \otimes e_{a})\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}[(\mathbf{F}^{T} \otimes \mathbf{I}_{m})vec(\mathbf{A}\mathbf{X}\mathbf{F}^{T}) + (\overline{\mathbf{F}}^{T} \otimes \mathbf{I}_{m})vec(\mathbf{A}\mathbf{X}\overline{\mathbf{F}}^{T})|vec(\mathbf{Y})] \\ &= (w^{T} \otimes e_{a})\left((\mathbf{F}^{T} \otimes \mathbf{I}_{m})vec(\mathbf{Y}) + (\overline{\mathbf{F}}^{T} \otimes \mathbf{I}_{m})\mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}}[vec(\mathbf{A}\mathbf{X}\overline{\mathbf{F}}^{T})|vec(\mathbf{Y})]\right) \\ &= (w^{T} \otimes e_{a})\left((\mathbf{F}^{T} \otimes \mathbf{I}_{m})vec(\mathbf{Y}) + (\overline{\mathbf{F}}^{T} \otimes \mathbf{I}_{m})vec\left(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T} + (\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^{T}\right)\right) \\ &= vec\left(e_{a}^{T}\left[\mathbf{Y}\mathbf{F} + \left(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T} + (\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^{T}\right)\overline{\mathbf{F}}\right]w\right) \\ &= e_{a}^{T}\left[\mathbf{Y}\mathbf{F} + \left(\mathbf{A}\mathbf{M}\overline{\mathbf{F}}^{T} + (\mathbf{Y} - \mathbf{A}\mathbf{M}\mathbf{F}^{T})(\mathbf{F}\mathbf{V}\mathbf{F}^{T})^{-1}\mathbf{F}\mathbf{V}\overline{\mathbf{F}}^{T}\right)\overline{\mathbf{F}}\right]w$$

where the sixth equality follows by Lemma 8.

B.9 Proof of Proposition 1

Note that

$$\begin{split} & K_{\theta}(P^*; \mathbf{A}, \mathbf{F}, \mu) \\ &= \kappa I_P(a; \mathbf{Y}) \\ &= \kappa \int_{\mathbf{Y} \in \mathcal{Y}} \sum_{a \in A(\mathbf{A})} P(a|\mathbf{Y}) \ln\left(\frac{P(a|\mathbf{Y})}{P(a)}\right) d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) \\ &= \kappa \int_{\mathbf{Y} \in \mathcal{Y}} \sum_{a \in A(\mathbf{A})} \frac{z_{\theta}(a, \mathbf{Y})}{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} \ln\left(\frac{z_{\theta}(a, \mathbf{Y})}{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} \frac{|A(\mathbf{A})|}{1}\right) d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) \\ &= \int_{\mathbf{Y} \in \mathcal{Y}} \sum_{a \in A(\mathbf{A})} \frac{z_{\theta}(a, \mathbf{Y})}{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} \left[\kappa \ln\left(z_{\theta}(a, \mathbf{Y})\right) - \kappa \ln\left(\frac{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})}{|A(\mathbf{A})|}\right)\right] d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) \\ &= \int_{\mathbf{Y} \in \mathcal{Y}} \left[\left(\sum_{a \in A(\mathbf{A})} \frac{z_{\theta}(a, \mathbf{Y}) \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}} \left[u_{\theta}(a, \mathbf{AX}) | \mathbf{Y}, \mathbf{A}, \mathbf{F}\right]}{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})}\right) - \kappa \ln\left(\frac{\sum_{c \in A(\mathbf{A})} z_{\theta}(c, \mathbf{Y})}{|A(\mathbf{A})|}\right) \right] d\mu_{\mathbf{Y}}(\mathbf{Y}|\mathbf{A}, \mathbf{F}) \end{split}$$

where the second equality follows from the definition of mutual information as the expected KL divergence and the third equality from plugging in the optimal choice probabilities of equation 10.

The result follows from the definition of equation 5

$$\begin{aligned} U_{\alpha}^{v}(\mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) \\ &= \sum_{a \in \mathcal{A}} P^{*}(a | \mathbf{Y}; \theta, \mathbf{A}, \mathbf{F}, \mu) R_{\alpha}^{v}(a, \mathbf{X}, \theta, \mathbf{A}, \mathbf{F}; \mu) \\ &= \sum_{a \in \mathcal{A}} P^{*}(a | \mathbf{Y}; \theta, \mathbf{A}, \mathbf{F}, \mu) v(a, \mathbf{A}\mathbf{X}, \theta) - (1 - \alpha) K_{\theta}(P^{*}; \mathbf{A}, \mathbf{F}, \mu) \\ &= \sum_{a \in \mathcal{A}(\mathbf{A})} \frac{z_{\theta}(a, \mathbf{Y}) u_{\theta}(a, \mathbf{X})}{\sum_{c \in \mathcal{A}(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} \\ &- (1 - \alpha) \int_{\mathbf{Y} \in \mathcal{Y}} \left[\left(\sum_{a \in \mathcal{A}(\mathbf{A})} \frac{z_{\theta}(a, \mathbf{Y}) \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}} \left[u_{\theta}(a, \mathbf{A}\mathbf{X}) | \mathbf{Y}, \mathbf{A}, \mathbf{F} \right]}{\sum_{c \in \mathcal{A}(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} \right] - \kappa \ln \left(\frac{\sum_{c \in \mathcal{A}(\mathbf{A})} z_{\theta}(c, \mathbf{Y})}{|\mathcal{A}(\mathbf{A})|} \right) \right] d\mu_{\mathbf{Y}}(\mathbf{Y} | \mathbf{A}, \mathbf{F}) \\ &= \sum_{a \in \mathcal{A}(\mathbf{A})} \left[\frac{z_{\theta}(a, \mathbf{Y}) u_{\theta}(a, \mathbf{A}\mathbf{X})}{\sum_{c \in \mathcal{A}(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} - (1 - \alpha) \int_{\mathbf{Y} \in \mathcal{Y}} \frac{z_{\theta}(a, \mathbf{Y}) \mathbb{E}_{\mu_{\mathbf{X}|\mathbf{Y}}} \left[u_{\theta}(a, \mathbf{A}\mathbf{X}) | \mathbf{Y}, \mathbf{A}, \mathbf{F} \right]}{\sum_{c \in \mathcal{A}(\mathbf{A})} z_{\theta}(c, \mathbf{Y})} d\mu_{\mathbf{Y}}(\mathbf{Y} | \mathbf{A}, \mathbf{F}) \right] \\ &+ (1 - \alpha) \kappa \int_{\mathbf{Y} \in \mathcal{Y}} \ln \left(\frac{1}{|\mathcal{A}(\mathbf{A})|} \sum_{c \in \mathcal{A}(\mathbf{A})} z_{\theta}(c, \mathbf{Y}) \right) d\mu_{\mathbf{Y}}(\mathbf{Y} | \mathbf{A}, \mathbf{F}) \tag{14}$$

C Appendix: Further Results for the Gaussian Case

C.1 Distribution of Vectorized Matrix Gaussian State

It may be helpful for better intuition to note that with the matrix Gaussian distribution of equation 12, the corresponding vectorized state is distributed according to

$$vec(\mathbf{X}^T) \sim \mathcal{N}(vec(\mathbf{M}^T), \mathbf{U} \otimes \mathbf{V}) = \mathcal{N}\left(\begin{pmatrix} \boldsymbol{m} \\ \vdots \\ \boldsymbol{m} \end{pmatrix}, \begin{pmatrix} \nu \mathbf{V} \quad c \mathbf{V} \quad c \mathbf{V} \\ c \mathbf{V} \quad \ddots \quad c \mathbf{V} \\ c \mathbf{V} \quad c \mathbf{V} \quad \nu \mathbf{V} \end{pmatrix}\right).$$

The vectorized filtered state is distributed according to

$$vec(\mathbf{Y}^T) = vec(\mathbf{F}\mathbf{X}^T\mathbf{A}^T) \sim \mathcal{N}(vec(\mathbf{F}\mathbf{M}^T\mathbf{A}^T), \mathbf{A}\mathbf{U}\mathbf{A}^T \otimes \mathbf{F}\mathbf{V}\mathbf{F}^T).$$

C.2 Joint Distribution of Linear Valuations of Options

Given the matrix Gaussian distribution of equation 12, the joint distribution of all options' values under linear utility is given by

$$\mathbf{X}\boldsymbol{w} \sim \mathcal{N}\left(\mathbf{M}\boldsymbol{w}, \mathbf{U}\boldsymbol{w}^T\mathbf{V}\boldsymbol{w}\right).$$
 (27)

C.3 Options' Joint Distribution of Attribute Values

Given diagonal $\mathbf{U} = \nu \mathbf{I}$, the joint distribution of each option's attribute values is, for all *a* identically and independently, given by

$$oldsymbol{e}_{a}^{T}\mathbf{X}\sim\mathcal{N}\left(oldsymbol{m},oldsymbol{e}_{a}^{T}\mathbf{U}oldsymbol{e}_{a}\mathbf{V}
ight)=\mathcal{N}\left(oldsymbol{m},
u\mathbf{V}
ight).$$

D Appendix of Section 3

Applying these general equations to the example of section 3, we obtain the distribution of the attributes' difference across choice options,

$$\begin{bmatrix} X_{\tilde{a}1} - X_{\tilde{b}1} \\ X_{\tilde{a}2} - X_{\tilde{b}2} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, 2(\nu - c) \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \right)$$

as well as, by equation 27, the distribution of both choice options' linear valuations

$$\begin{bmatrix} w_1 X_{\tilde{a}1} + w_2 X_{\tilde{a}2} \\ w_1 X_{\tilde{b}1} + w_2 X_{\tilde{b}2} \end{bmatrix} \sim \mathcal{N} \left(\mathbf{0}, (w_1^2 + 2\rho w_1 w_2 + w_2^2) \begin{bmatrix} \nu & c \\ c & \nu \end{bmatrix} \right),$$

both of which are plotted in Figure 6.



Figure 6: The plots in the left column show the joint distribution of $(X_{\tilde{a}1} - X_{\tilde{b}1}, X_{\tilde{a}2} - X_{\tilde{b}2})$ under the prior, each row for a different set of covariance parameters. The plots in the right three columns show the joint distribution of $(w_1X_{\tilde{a}1} + w_2X_{\tilde{a}2}, w_1X_{\tilde{b}1} + w_2X_{\tilde{b}2})$, each column for a different \boldsymbol{w} . The figure's left-most column illustrates how the attribute covariance ρ rotates the distribution, while the action covariance c affects the distribution's dispersion. The remaining columns illustrate how the joint distribution of the two option's utility has a dispersion that depends on the interplay of preference weights \boldsymbol{w} and attribute covariance ρ , while the covariance is determined by the action covariance c.



Figure 7: Gross welfare $U_1^u(\mathbf{X}, \theta, \mathbf{A} = \mathbf{I}, \mathbf{F}; \mu)$ for the three different information policies \mathbf{F} , now shown as a function of the preference weights $\boldsymbol{w} = (w_1, w_2)$. Each row corresponds to a given state \mathbf{X} and covariance ρ .