

Information Theory and Observational Limitations in Decision Making

David Wolpert
David Leslie

SFI WORKING PAPER: 2011-09-045

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

Information Theory and Observational Limitations in Decision Making

David H. Wolpert

Santa Fe Institute, 1399 Hyde Park Road Santa Fe, NM 87501

Center for Nonlinear Studies, MS B258, Los Alamos National Laboratory, Los Alamos, NM 87545

MS 269-1, NASA Ames Research Center, Moffett Field, CA, 94035, USA

david.h.wolpert@nasa.gov

David S. Leslie

School of Mathematics, University of Bristol,

University Walk, Bristol, BS8 1TW, UK

david.leslie@bristol.ac.uk

September 9, 2011

Abstract

We introduce a general framework for formalizing and analyzing the problem faced by a Decision Maker (DM) working under information-theoretic constraints on their observational ability. The random utility model and the “hedonic utility” model of Netzer and Robson (NR) are special cases of this framework. We begin by applying information theory to our framework to derive general results concerning the expected regret of DM under observational limitations. We then turn our attention to the effects of observational limitations on choice behavior (rather than the regret values induced by that behavior), focusing on the special case of NR. First we provide a simple derivation of two assumptions made by NR, and then of the result of NR that a particular hedonic utility function satisfies certain optimality principles. We then extend NR to allow a countable rather than uncountable set of states of the world. In particular we show how to use dynamic programming to solve for the optimal preference order of DM in this extension. We also extend NR by considering the case where more than two options are presented to DM , showing that the results of NR change in such a case.

Keywords: Hedonic utility, risk attitudes, S-shaped utility, decision theory, constrained observation, bounded rationality, information theory, regret

JEL classification: D01, D81, D87

1 Introduction

One of the enduring problems of economics is explaining the apparent non-rationality of human decision behaviour [3, 43, 4, 22], both in strategic and non-strategic settings. Several explanations have been explored in a strategic setting. For example, a large body of work explains systematic non-rationality in any *single* instance of a game by considering what is rational behaviour for an infinite sequence of repetitions of that game. Early examples of this approach are reciprocal altruism and the folk theorems [31, 24, 12]. Other explanations also consider an infinite sequence of repetitions of the game, but analyze an evolutionary dynamics of the players' strategies across that sequence [3, 31, 11, 17, 46, 2, 17, 14, 33, 9, 23, 30, 7]. A related body of work has allowed the preferences of the players to evolve instead of their strategies [19, 18, 38, 6, 15, 1].

Some other work instead addresses the strategic setting by assuming that there are some cognitive limitations of the players, and then examines what optimal behaviour is subject to those limitations. Early examples of this work have considered limitations in the computational power of the players, by explicitly modelling their computational algorithm [11, 16, 36, 37, 13].

More recently, some work has concentrated on how cognitive limitations might explain non-rationality in non-strategic situations, i.e., in situations where a single person must choose among a set of options. Typically, this work does not consider limitations in the computational abilities of the Decision Maker (DM). Rather it considers limitations in the ability of DM to observe all characteristics of their various options with perfect accuracy and the implications of those limitations.

A prominent early example of this approach is the Random Utility Model (RUM) of conventional choice analysis [44, 29, 45, 28]. Randomness in RUM is sometimes considered as taking place at the demographic level (e.g. [29]). In other

interpretations of the model (e.g. [44]) one assumes that DM is simultaneously offered all options y in some set Y , but forms noisy estimates / observations of the associated utilities, $\psi(y)$. One then assumes that DM will choose whichever of their options they observe to have the highest utility value, without concern for the noise in their observations of those utility values.

More recent work on non-rationality in choice differs in two major respects from RUM. First, the observational limitations of DM in observing utilities $\psi(y)$ are not modelled as random noise corrupting DM's observations of those utility values. Rather the observational limitations are modelled by having DM observe the value of a non-invertible, parameterised "hedonic utility function", $V(y) \in \mathbb{R}$. In these models, DM must choose among a set of options based on the value of $V(y)$, whereas what they want to maximise is $\psi(y)$. Often $V(y)$ is constrained to be a single-valued function of $\psi(y)$, e.g., a discretization of $\psi(y)$. To emphasize the fact that $\psi(y)$ is what DM should optimize according to the axioms of decision theory, we will refer to it as the **Savage** utility function, in contrast to the hedonic utility function.¹

Second, in this recent work, while there is stochasticity in the decision problem, it does not arise in the observations made by DM concerning the utility values of their options, as it does in some interpretations of RUM. Rather in this recent work, randomness arises because the set of options to be considered is not all of Y (as in RUM), but is instead a randomly generated subset of

¹In some of the work discussed below, this function is referred to as a "fitness" function. However we do not want to limit ourselves to an evolutionary context.

Y [35, 10, 34, 32, 25].

In this paper we introduce a general framework for modelling observational limitations of decision makers, which includes both RUM and the more recent work as special cases. Our goal is to understand the implications of the unavoidable inability of real world decision makers to perfectly observe all aspects of an option y that are relevant to the value of $\psi(y)$. Exploiting the insights of information theory [5, 27], we quantify such a limited observation process with an information channel P that stochastically maps a set of k options from Y together with the associated k values $\psi(y)$ to some abstract space of k observation values, $\mathcal{T}^k = \times_{i=1}^k \mathcal{T}$.

Often the information channel will apply the same stochastic mapping to each of the k separate pairs $(y, \psi(y))$ to produce an associated value t . As an important example of such an information channel, \mathcal{T}^k could provide a discretized version of each of the k values of $\psi(y)$. This is what is done in the recent work mentioned above, e.g., in the model introduced by Netzer [32]. (This model is based on a model of Robson [35], and henceforth simply referred to as “NR”.)

As another, more concrete example, say that Y consists of all possible fruits, with all possible shapes, and all possible optical reflection spectra. Due to limitations of the human eye, people cannot discriminate all shapes to infinite accuracy, and can only register red, green and blue wavelengths (and even those only to limited precision). Moreover, the process of registering an image on the human retina is a noisy one. Accordingly, the production of an image on a human retina of a particular fruit $y \in Y$ can be modelled as the output of a stochastic information channel, $P(t | y)$ (independent of $\psi(y)$), where $t \in \mathcal{T}$ represents the signal sent from the retina down the optic nerve.² Note that there is information compression and unavoidable noise in going from y to the signal in the optic nerve. Note also that the space of such signals does not involve a discretization of the space of all possible fruit, as there is in the NR model. (See [40, 8] for some of the large literature on analyses of the neurobiological vision system as an information channel.)

Our model presumes the existence of such an information channel limiting how much DM knows about their options, but is otherwise very general. To be precise, we assume that for some $k \in \mathbb{N}$, there is a distribution G_k over Y^k that generates the sample of k elements of Y .³ DM must choose among the k elements of the sample, and wishes to make the choice y with maximal value of an associated Savage utility $\psi(y)$. However DM cannot directly observe which elements of Y are in the sample and/or their associated Savage utility values. Rather the elements of the sample and the associated utilities are stochastically mapped through the information channel $P(\vec{t} | \vec{y}, \vec{\psi})$ into the space \mathcal{T}^k . DM then applies a “hedonic utility function” $U : \mathcal{T} \rightarrow \mathbb{R}$ to rank each of the k components of \vec{t} , and chooses the one with the highest such rank, resolving ties arbitrarily.⁴

So in our framework the following sequence occurs: A distribution G_k is sampled to produce a random \vec{y} , and a probability density function is sampled to produce the associated Savage utility values $\vec{\psi}$. The pair $(\vec{y}, \vec{\psi})$ is then observed stochastically as \vec{t} . This is then used by DM together

²This example’s model of vision is only an approximation, since neural information is actually encoded in spike trains and chemical gradients.

³Throughout this paper we will use the term “distribution” to include cumulative distributions, probability distributions and probability density functions, relying on the context to make the precise meaning clear when that is important.

⁴Note that for simplicity we use the term “hedonic utility function”, even though unlike the function V considered in NR, U ’s domain is \mathcal{T} , not Y .

with their hedonic utility function U to make their choice. (A more complete statement of the framework is presented below.)

This is a very broad framework. In particular, the RUM model is a special case. To see this, note that in the RUM model, every element of a finite Y is considered by DM, without any stochasticity. Similarly, in the RUM model, the utility function $\psi(\cdot)$ is pre-fixed, without any stochasticity. Hence to cast RUM within our model we simply take $k = |Y|$ and set G_k to be the delta function on a k -vector consisting of one instance of each element of Y . In the RUM model, the feature of an option observed by DM is a real-valued utility, so \mathcal{T} is simply taken to be \mathbb{R} with the hedonic utility function $U : \mathcal{T} \rightarrow \mathbb{R}$ being the identity. Finally P is a channel stochastically mapping elements y to noisy estimates of the associated value $\psi(y)$.

As mentioned, NR is also a special case of our framework. In contrast with the RUM model, in NR only two options are considered, and these are sampled independently from Y . Therefore $k = 2$, and $G_2(y_1, y_2) = G(y_1)G(y_2)$ for some distribution G . Just like in RUM, $\psi(\cdot)$ is pre-fixed. There is a set of T bins of the set of all possible values, $\{\psi(y) : y \in Y\}$, which are labelled $t = 1, \dots, T$. (We adopt the convention that the label of a bin increases with the $\psi(y)$ values in the bin.) These bins comprise \mathcal{T} . The information channel maps $(y, \psi(y))$ to the associated bin containing $\psi(y)$. More formally, write the single-valued mapping that takes $\psi(y)$ to the associated bin as $b : \psi(Y) \rightarrow \mathcal{T}$. Then in NR, unlike in the RUM model, the information channel P is deterministic; it can be expressed as $P(t | \psi(y)) = \delta(t - b(\psi(y)))$. NR then takes $U(t) = t/T$, and the ‘‘hedonic utility’’ analysed by NR is the function over Y given by $V(y) = U(b(\psi(y)))$.

Our framework generalizes both RUM and NR in several respects. Firstly, in our framework we do not make any *a priori* assumptions about the spaces \mathcal{T} and Y . Secondly, we assume only that the limited ability of DM to map from $Y \times \psi(Y)$ to \mathcal{T} can be modeled as an information channel $P(\vec{t} | \vec{y}, \vec{\psi})$. In some of our analysis, we take the channel and U as given, and analyze the implications. (This is similar to the typical RUM analysis.) In most of our paper though, we assume that the information channel and U are optimized subject to some constraints. (This is similar to the analysis of NR.) The presumption behind such analysis is that a fully rational DM will select optimal P and U subject to these constraints, and the analysis concerns the implications of the constraints on their choice.⁵

This generality of the framework means that it applies to individuals facing many different types of decision problems, in which different constraints are placed on P and/or U . The presumption is that changing those constraints on P and/or U by changing the problem results in changing what P and/or U DM uses. Physically, that would mean that the relative hedonic utility assigned by DM to two options may change with the context in which those options are presented to DM.

One might expect that since any conditional distribution P is an information channel, the constraints on P might usefully be expressed in terms of information theoretic quantities, like the channel’s capacity, or the mutual information between inputs and outputs of the channel [5, 27]. This is indeed the case. As an example, in Sec. 3.2 we relate the expected regret (in the decision-theoretic sense of the term) of the choice made by DM to the amount of mutual information in the

⁵In this article we do not address the question of how this optimisation may be carried out, noting only that while evolution (as considered by NR) may be one mechanism, a flexible lifetime learning mechanism by DM may also result in optimal P and U .

channel. Since in the real world DM can often pay for extra mutual information (e.g., by purchasing more fiber optic), this result can be used to determine how much DM should pay to upgrade the channel. As another example, a standard result of information theory gives us a law of diminishing returns concerning how much DM should pay to improve the information capacity.

After illustrating the economic consequences of such information theoretic results, we demonstrate the implications of our model within the NR framework, in which the only constraint on P we consider is that its range must be the fixed space \mathcal{T} .

Of course, information theory has been used in economics before our work, e.g., in the finance and macroeconomics literature. In fact, it has even been used before to analyze seemingly anomalous human choice behavior in providing an explanation for “rational inattention” [41, 42], the phenomenon by which decision makers sometimes elect not to expend the effort to acquire all information that may be relevant to a decision.

Our precise contributions, in the order that we make them, are as follows:

1. We introduce a fully formal version of the framework described above for modelling the limitations of human decision makers arising from constrained information channels (Section 2).

We elaborate the relationship between constraints on the information channel used by DM and the expected regret that DM has over decisions they make based on the information in that channel. In particular we relate the amount that DM pays to get a channel of a given fidelity and the resultant regret they experience over decisions based on that channel. (Section 3).

For the remainder of the paper we turn our attention to the effects of observational limitations on choice behavior rather than the utility values induced by those effects. We focus on the special case of NR.

2. We show that when constraints are placed on the observation space \mathcal{T} instead of on the information channel, an optimality criterion recovers some of NR’s assumptions result. In particular if $k = 2$ and \mathcal{T} is finite then a channel P corresponding to a deterministic discretisation of the space $\psi(Y)$ is optimal (Section 4.1).
3. For the precise case that NR considers, where \mathcal{T} is finite, $\psi(Y)$ is a subinterval of the reals, and $k = 2$, we derive the “centre of mass rule” relating the optimal decision boundaries (Section 4.3). We use this to directly establish how an S -shaped hedonic utility function can be optimal: that shape arises from the combination of a clustering of decision boundaries with an assumption that the jumps in hedonic utility as one crosses those boundaries be constant (Section 4.4).⁶
4. We expand this analysis to the case where $\psi(Y)$ is finite, with \mathcal{T} also finite and $k = 2$. In many situations, especially those studied in experiments, the set of available options is

⁶It is important to note that the constant-jump assumption is arbitrary, and when it is dropped a non- S -shaped hedonic utility can trivially result. Accordingly these analyses based on decision boundaries (NR’s and that in the latter part of our article) cannot, by themselves, explain the experimental observations of S -shaped utility functions.

indeed finite, and therefore a model which relies a continuum $\psi(Y)$ of possible utility values is troublesome. We show how to use dynamic programming to solve for optimal decision boundaries in this case (Section 4.5).

5. We extend NR to the case where $k = 3$ options are simultaneously presented to DM (Section 5). We show that the optimal hedonic utility function may be different for this case from when $k = 2$ options are presented. Accordingly, if DM knows k , then they should vary the hedonic utility function they use to choose among their options depending on that value of k . In other words, in loose agreement with experiment, DM should vary the relative value they assign to different options depending how many options they have. On the other hand, say that DM does not know the value of k . In this case, the best DM can do is assume some distribution $P(k)$ of decision problems they encounter in their lifetimes, and solve for the (single) hedonic utility function that optimizes expected Savage utility under $P(k)$.

We end with a conclusion of our results. Our notation, and some technical proofs, are placed in appendices.

2 Framework formalization

In this section we present our information channel framework in full.

We assume individuals are faced with a choice between k options $y_1, \dots, y_k \in Y$. That set is formed by sampling a distribution G_k . There are no *a priori* restrictions on Y or G . We write $y_{1:k} \triangleq (y_1, \dots, y_k)$, and when k is obvious we sometimes abbreviate this as \vec{y} .

In addition to these options there is a function $\psi : Y \rightarrow \mathbb{R}$ giving the Savage utility of any element in Y . We will denote by $\psi(Y)$ the set of all possible Savage utility values. Note that for finite Y , any function $\psi(\cdot)$ is an element of $\psi(Y)^{|Y|}$. We put no restrictions on the function $\psi(\cdot)$. In particular, we will sometimes treat it as a random variable, and we will sometimes take it to be pre-fixed in the problem definition.⁷

We assume that before choosing among any set of k options \vec{y} , observational limitations force DM to apply an information channel P that takes \vec{y} and the associated k -vector $\vec{\psi} = (\psi(y_1), \psi(y_2), \dots, \psi(y_k))$ into a product space $\mathcal{T}^k = \times_{i=1}^k \mathcal{T}$. Formally, with $\vec{y}, \vec{\psi}$ and $\vec{t} \in \mathcal{T}^k$ being random variables, P specifies the conditional probability distribution $P(\vec{t} | \vec{y}, \vec{\psi})$. DM then chooses among their k options by considering the associated components of \vec{t} . To do this, DM uses a ‘‘hedonic utility function’’ $U : \mathcal{T} \rightarrow \mathbb{R}$; DM chooses the option given by $\operatorname{argmax}_{i=1, \dots, k} U(t_i)$, choosing randomly among the maximizers if there is more than one.

Note that there are two utility spaces involved in this model: $\psi(Y)$ (‘‘Savage’’ utility) and $U(\mathcal{T})$ (‘‘hedonic’’ utility). In this sense the framework is related to the evolution of preferences literature [19, 18, 38, 6, 15, 1], and to recent work on persona games [21, 20].

⁷Formally, we treat the latter as a special case of the former. For example, for finite Y , we view a ‘‘pre-fixed’’ ψ as meaning ψ is generated from a probability density function over $\psi(Y)^{|Y|}$ that is a Dirac delta function.

3 Information-theoretic results

One of the advantages of our casting decision problems in terms of information channels is that by exploiting the theorems of information theory we can relate information-theoretic characteristics of a choice scenario and the associated expected utility of DM. In this section we present three such relations, all quite simple, illustrating the breadth of our framework as we do so.

These relations between information-theoretic characteristics of a scenario and the associated expected utility of DM are loosely analogous to relations between the characteristics of the additive random utility in an RUM and the associated expected utility of DM. Note in particular that additive random utility is a non-physical quantity, so experiment cannot directly measure it. Rather experimental data can only infer a random utility indirectly by examining the choices ultimately made by DM. The situation is not quite so challenging for analyzing choices based on noisy information channels. In addition to examining the choices made by DM, we can sometimes investigate the information channels underlying those choices directly, since information channels are conditional probability distributions, and therefore often instantiated in physically observable ways.

3.1 Background

We start by presenting some basic information theory. For any random variable A taking values a with probability $P(a)$, the **Shannon entropy** of that variable is $H(A) \triangleq -\sum_a P(a)\ln[P(a)]$.⁸ The entropy of a random variable is largest when the associated distribution is uniform, and it equals 0 if the associated distribution is a delta function. So entropy measures “how spread out” a distribution is. Alternatively, it can be viewed as telling you how much information is in that distribution; a delta function $P(a)$ (zero entropy) contains maximal information about a , whereas a uniform $P(a)$ (maximal entropy) contains no information about a .

Similarly, if we also have a random variable B taking values b , then the **conditional entropy** of A given B is

$$\begin{aligned} H(A | B) &\triangleq -\sum_b P(b) \sum_a P(a | b)\ln[P(a | b)] \\ &= -\sum_{a,b} P(a, b)\ln[P(a | b)] \end{aligned} \tag{1}$$

(Unless explicitly stated otherwise, all logarithms in entropies are assumed to be base 2.) This is a measure of how much knowing b tells you about what a is, averaged according to the probability of the various b 's. Note that

$$H(A | B) = H(A, B) - H(B)$$

where $H(A, B) = -\sum_{a,b} P(a, b)\ln[P(a, b)]$.

⁸An integral replace the sum if a is uncountable, and a measure is introduced into the ln. See [5, 27] for details.

The **mutual information** between A and B is a measure of the “co-variation” of the random variables A and B . It is defined as

$$\begin{aligned} I(A; B) &= H(A) - H(A | B) \\ &= H(A) + H(B) - H(A, B) \\ &= I(B; A). \end{aligned} \tag{2}$$

The mutual information between A and B quantifies how much extra knowing the channel output b tells you about the channel input a , beyond what the distribution $P(a)$ tells you. Due to its symmetry, that mutual information also tells you how extra much knowing the channel input a tells you about the channel output b , beyond what the distribution $P(b)$ tells you. If A and B are statistically independent, their mutual information equals 0 (the smallest possible value). In contrast, if knowing b fixes a exactly (and vice-versa), the mutual information is maximal.

Intuitively, mutual information can be thought of as an information-theoretic analog of the correlation of real-valued random variables. However it has several substantial advantages over the correlation for quantifying how much two variables co-vary. Most obviously, the mutual information of two variables is perfectly well-defined even if the variables are symbolic-valued rather than real-valued. In addition, mutual information is not subject to the artifacts that plague correlation. As an example, let A and B be real-valued random variables, where $P(a, b) \propto \delta(a^2 + b^2 - 1)$ where $\delta(\cdot)$ is the Dirac delta function. Then the correlation of those two variables is zero, even though knowing the value of a tells you a huge amount concerning the value of b (and vice-versa). In contrast, their mutual information is close to maximal, reflecting the tight relation of the two random variables.⁹ For these reasons, mutual information is a natural choice for how to measure how closely two random variables are related.

Finally, a key concept of modern telecommunication theory is the **information capacity** of a channel, which is a measure of how much information can be transmitted down a channel [5, 27]. The information capacity is defined to be the maximal mutual information between the input to the channel and the output of the channel, with maximization taken over the distribution of the channel input. Formally, if the channel input is the random variable A taking values a , and the output is the random variable B taking values b , then the channel capacity is $\sup_{P(A)} I(A, B)$.

Using these definitions, standard results from information theory can be applied to our framework, giving insights into the types of decision making we expect to observe. The rest of this section illustrates this with several examples.

3.2 Implications of rate distortion theory

We start by considering a general scenario, where there are no restrictions on G_k . It could be a distribution where the components of \vec{y} are all statistically coupled, as in RUM, or one where they are formed by IID sampling a distribution over Y , like in NR, or something else entirely. Similarly,

⁹Strictly speaking, to make the definition of mutual information for this scenario precise, one has to choose a measure over \mathbb{R}^2 . Or more simply, one could simply restrict the values of a and b to lie in a fine-grained discretization of \mathbb{R}^2 , so that the definition of mutual information involves sums rather than integrals.

we put no restrictions on $P(\vec{t} | \vec{y}, \vec{\psi})$. The only thing we rely on is that $P(\psi)$ is not a Dirac delta function.

Some bounds that apply to all such scenarios arise from a branch of information theory called “rate distortion theory”, which provides relations between the expected value of a cost function and the information capacity of a channel. To begin, suppose there are two sets A and B , and a cost function $d : (a \in A, b \in B) \rightarrow \mathbb{R}$ that is non-negative and equal to 0 if $a = b$. There is also a fixed prior $P(a)$. The goal is to construct a channel $P(b | a)$ that will minimize expected cost.

Physically, that channel may involve observational devices, computational devices, amounts of fiber optic cable, etc. Unfortunately, such channels cannot be arbitrarily accurate, since accuracy is expensive. Accordingly, suppose there is a fixed budget that can be spent by DM on making the channel accurate. For example, we can imagine that DM must pay money to a provider of fiber optic, and the provider charges DM more money for bigger bandwidth cables that will be more accurate information channels for any given input stream. (The funds that DM gives to such a provider are not reflected in the cost d ; the latter is a function whose value depends on the relation between an input to the channel and the output.)

Since $P(a)$ is fixed and the goal is to set $P(b | a)$, one way to quantify the constraint on how “accurate” the channel is as an upper bound on the mutual information $I(A; B)$ resulting from the choice of $P(b | a)$. If the bound is I^* — if for example due to budgetary constraints DM cannot afford to pay for a channel that is more accurate than that — then the best value for the associated expected cost that can be attained is defined as the **distortion rate function**,

$$D(I^*) \triangleq \min_{P(b|a): I(A;B) \leq I^*} [\mathbb{E}(d)] \quad (3)$$

where the expectation of d is evaluated under the distribution $P(a)P(b | a)$.

Conversely, suppose the expected cost must be bounded above, say by d^* , and the goal is to find the channel with smallest mutual information between A and B that achieves this expected cost. That smallest mutual information is given by the **rate distortion function**,

$$R(d^*) \triangleq \min_{P(b|a): \mathbb{E}(d) \leq d^*} [I(A; B)] \quad (4)$$

The rate distortion function is the inverse of the distortion rate function, i.e., $R(D(I^*)) = I^*$ for all I^* . In addition, both functions are decreasing and everywhere convex [5, 27].

To apply these results to the economic framework, say we are given a particular hedonic utility $U(\cdot)$ that will be used to map any \vec{t} to a choice. Then the **regret** of DM (in the standard, economic sense) for any particular $\vec{\psi}$ and \vec{t} is

$$C(\vec{\psi}; \vec{t}) = \max_j [\psi_j] - \psi_{\text{argmax}_i [U(t_i)]} \quad (5)$$

where ψ_i is shorthand for the i 'th component of $\vec{\psi}$. Note that the regret function is non-negative and equal to 0 if $\vec{\psi} = \vec{t}$. Accordingly, we can view it as a cost function in the sense of rate distortion theory, for a channel with input $\vec{\psi}$ and output \vec{t} .

Due to this, the results from distortion rate theory discussed above mean that both the rate distortion function and the distortion rate function are convex and decreasing if we use regret for the cost. This allows us to formally establish some intuitively reasonable properties of many

choice scenarios. In particular, it means that for any scenario describable with our framework, the expected regret decreases as we increase the budget for channel accuracy. Less trivially, the convexity of the distortion rate function means that a law of diminishing returns applies to the extra effort / expenditures involved in improving the channel accuracy:

Proposition 1. *Fix the input distribution $P(\vec{y}, \vec{\psi})$. Any channel $P^\dagger(\vec{t} | \vec{y}, \vec{\psi})$ results in an associated mutual information $I(P^\dagger)$ between \vec{t} and $(\vec{y}, \vec{\psi})$. Define $\mathcal{M}(I^*)$ to be the multifunction taking a mutual information value I^* to the set of prices that DM would have to pay to use any of those channels P^\dagger that have mutual information $I(P^\dagger) = I^*$. Assume that $\mathcal{M}(I^*)$ is a single-valued, convex, differentiable function of I^* with positive-definite derivative. Then the minimal expected regret of DM for using a channel that has a price π is a decreasing convex function of π . (The minimization is over all channels with price π and over all ways that DM might use the channel.)*

Proof. If DM has a budget of π , then the minimal expected regret of DM is $D[\mathcal{M}^{-1}(\pi)]$. Since \mathcal{M} has positive-definite derivative and is convex, its inverse has positive-definite derivative and is concave, i.e., $\mathcal{M}^{-1}(\pi)$ is an increasing, concave function of π . This implies two things. First, since $D(\cdot)$ is a decreasing function, it means that $D[\mathcal{M}^{-1}(\cdot)]$ is a decreasing function. Second, since $D(\cdot)$ is in fact a decreasing convex function, it means that the double-derivative of $D[\mathcal{M}^{-1}(\pi)]$ with respect to π is non-negative. So $D[\mathcal{M}^{-1}(\pi)]$ is a convex function of π . \square

It is possible that the assumption in Prop. 1 about the single-valuedness of \mathcal{M} does not hold. However it is often a reasonable assumption. In particular, often hardware limitations will mean that the cost to construct any channel $P(\vec{t} | \vec{y}, \vec{\psi})$ can be approximated as a single-valued function of the associated mutual information $I(\vec{t}; \vec{y}, \vec{\psi})$. The assumption in Prop. 1 concerning the form of \mathcal{M} should hold even more often; it amounts to the standard assumption concerning the shape of an economic supply curve. For the rest of this subsection we will implicitly take these assumptions to hold.

DM wants their expected regret to be small, and also wants the price π they pay for using their channel to be small. For simplicity, formalize this by saying that the preference function of DM in choosing how much to pay is a linear combination of their expected regret and how much they pay. Assume also that if they pay for a channel with mutual information I^* , then DM can choose the actual such channel $P^\dagger(\vec{t} | \vec{y}, \vec{\psi})$ that they get. So assuming they are rational, for any given amount of mutual information they pay for, I^* , the associated channel $P(\vec{t} | \vec{\psi}, \vec{y})$ they get is the best possible one for them of all channels that have $I(\vec{t}; \vec{\psi}, \vec{y}) \leq I^*$, in the sense that they cannot get lower expected regret using any of those other channels than using the channel they choose. Under these assumptions, Prop. 1 means that (under its assumptions concerning the price function \mathcal{M}) the preference function is concave with a single maximum; there is a unique price that a rational DM will pay to maximize their preference function.

Example 1. *Suppose that DM is a politician in office, who has to choose one of k mutually contradictory bills to vote for. Those k bills comprise Y . For simplicity we assume they will vote for one and only one of the bills. Suppose further that all the politician cares about is being re-elected, and $\psi(y)$ is the probability that they get re-elected if they vote for bill y . Hence their regret is just*

the maximal re-election probability over all bills they might vote for, minus the re-election probability given the bill they actually do vote for. Their expected regret is simply the expectation of this difference, where the expectation is over their possible choices of what bill to vote for.

To have as much information about how to vote as possible, DM hires a company to poll their electorate. The polling company will return to DM an estimate of the probability of their being re-elected for each of the bills they might vote for. That vector of estimates is \vec{t} ; the polling company is the information channel $P(\vec{t} | \vec{\psi}, \vec{y})$. (Note that as in an RUM, G_k is a delta function specifying that every $y \in Y$ is sent through the information channel.) For simplicity take $U(t) = t$ for all t , so by choosing to vote for the bill i that maximizes $U(t_i)$, DM will be voting for the bill that the polling company says is most likely to get them re-elected. Perfectly accurate estimates \vec{t} of $\vec{\psi}$ are unachievable. Accordingly the bill y' that DM chooses based on the polling company's estimated probabilities in general will not be the same as the bill y that in fact would have maximized their chances of getting re-elected had they voted for it. That drop in their probability of being re-elected is the regret of DM.

Assume that the polling company can provide improved estimates \vec{t} , but at extra cost to the polling company, and therefore to DM. So if DM pays more, the information channel will be more accurate. Model this by having there be a single-valued convex non-decreasing function between the fee that DM pays to the polling company and the mutual information of the results of the poll, $I(\vec{t}; \vec{\psi}, \vec{y})$.

Assume that the goal of DM in deciding how much money to give the polling company is minimizing a sum of their expected regret and the money they pay. Prop. 1 then tells us that there is only one possible amount that a rational DM will pay to the polling company.

For special cases we can go further than Prop. 1's general convexity result, and evaluate the rate distortion and distortion rate functions exactly. An example is given by the following result:

Proposition 2. Say that $k = |Y|$ and G_k always produces a \vec{y} running over all Y , as in a RUM. Suppose each successive component of $\vec{\psi}$ is generated by sampling the same Gaussian distribution that has standard deviation σ . Then if the mutual information between \vec{t} and $\vec{\psi}$ is no more than I^* , the minimal expected regret is $\frac{\sigma}{\sqrt{2}} e^{-I^*/k}$. Conversely, to guarantee that the expected regret is no worse than c^* , DM must be willing to pay enough to ensure that the mutual information is at least $k \ln[\frac{\sigma}{c^* \sqrt{2}}]$.

Proof. To begin, choose a different cost function, $d(\vec{\psi}, \vec{t}) = |\vec{\psi} - \vec{t}|^2$. Then it can be shown [5] that

$$R(d^*) = \frac{k}{2} \ln\left(\frac{\sigma^2}{d^*}\right).$$

Inverting,

$$D(I^*) = \sigma^2 e^{-2I^*/k}.$$

We must convert this bound on expected d into a bound on expected regret. To begin, define i^* as the index of the component of $\vec{\psi}$ that is largest, and define j^* as the index of the component of \vec{t} that is largest. So the regret is given by

$$C(\vec{t}, \vec{\psi}) = \psi_{i^*} - \psi_{j^*}.$$

Next, as shorthand, define $\omega \equiv \sqrt{[\vec{\psi} - \vec{t}]^2}$. It is straightforward to verify that given any value of $C(\vec{t}, \vec{\psi})$, the smallest that ω^2 can be is $2[C(\vec{t}, \vec{\psi})]^2$. (This limit occurs if \vec{t} and $\vec{\psi}$ are identical in all components except t^* and j^* .) So given a value of ω , the largest $C(\vec{t}, \vec{\psi})$ can be is $\omega/\sqrt{2}$. Moreover, given a value $V \equiv \mathbb{E}(\omega^2)$, $\mathbb{E}(\omega) \leq \sqrt{V}$. Combining, we see that $\mathbb{E}(C) \leq \sqrt{V}/2$. In turn, V is given by the formula above for $D(I^*)$. Plugging that in completes the proof. \square

By the Shannon lower bound theorem, if $P(\vec{\psi})$ has standard deviation σ but is not a Gaussian, then the distortion rate function is worse than it is for a Gaussian $P(\vec{\psi})$ with standard deviation σ . That means that for non-Gaussian $P(\vec{\psi})$ with standard deviation σ , DM must pay more (in terms of the mutual information they purchase) to guarantee the same expected regret as they would for a Gaussian $P(\vec{\psi})$ with standard deviation σ .

Example 2. Return to Ex. 1. Assume that whatever bill DM votes for will not have a very large effect on the re-election chances of DM, which are likely to be close to 50% for all of their potential votes. So for any particular bill y , the density function $P(\psi(y))$ looks approximately like a Gaussian truncated to lie in $[0, 1]$. Also assume for simplicity that the re-election chances of DM conditioned on his voting for bill y and their chances conditioned on voting for bill y' are statistically independent for any $y, y' \neq y$. More formally, assume that to a good approximation, each component of $\vec{\psi}$ is generated by independently sampling the same distribution, which has a standard deviation σ which is far smaller than 1, the width of possible values of any $\psi(y)$.

Given this, Prop. 2 tells us that the expected regret of DM is at least $\frac{\sigma}{\sqrt{2}}e^{-I^*/k}$. (Recall from the remark following Prop. 2 the regret will be bounded below by the regret that would be achieved if the distributions were exactly Gaussian.) Note the exponential decrease in expected regret, so that once I^* is large, the benefit of increasing it further is negligible. Note also that the smaller the value of k , the number of bills DM has to choose among, the better it is for DM.

Say we are given the function taking a given mutual information I between re-election probability predictions by the polling company and the actual re-election probabilities, and the price that the polling company charges for a poll with mutual information I . We can plug in that function to rewrite our lower bound on the expected regret of DM as a function of the price they pay to the polling company. Of course, polling companies do not typically quote prices in terms of mutual information. So to evaluate this bound on real-world data, some inference and analysis would be necessary to translate the quantities used by polling companies to set prices (number of people they poll, how many questions, etc.) to values of mutual information.

3.3 Implications of Fano's inequality

Assume as in Prop. 2 that DM receives information about all of their options. Suppose that in addition, the utility of DM is 1 if they make a ‘‘correct’’ choice, and zero otherwise. Such a situation would exist if the problem was to search in a number of locations for an item of food, for example. For this situation we can extend the results of the previous section.

Proposition 3. Let $k = |Y|$, and let G_k be a delta function specifying that all y 's are sent through the data channel, as in a RUM. Have every $\vec{\psi}$ specify some particular associated y value such that

the utility of DM equals 1 if they pick that y value, and 0 otherwise. So the regret $\epsilon \equiv C(\vec{\psi}; \vec{t})$ is a random variable that equals either 0 or 1. Then

$$i) H(\epsilon) + P(\epsilon = 1)\ln_2[k] \geq H(\vec{\psi} | \vec{t})$$

$$ii) I(\vec{\psi}; \vec{t}) \geq H(\vec{\psi}) - H(\epsilon) - P(\epsilon = 1)\ln_2[k]$$

Proof. The proof of the first result is immediate by applying Fano's inequality [5, 27]. The second result follows by using the general fact that $I(A; B) = H(A) - H(A | B)$ for any pair of random variables A and B , applied with $A = \vec{\psi}$ and $B = \vec{t}$. \square

If the expected regret $\mathbb{E}(\epsilon) = \mathbb{P}(\epsilon = 1)$ is near 0, then since ϵ is binary-valued, $H(\epsilon)$ is also necessarily near 0. So Prop. 3(ii) tells us that in the limit where the expected regret of DM approaches zero, the mutual information $I(\vec{\psi}; \vec{t})$ equals or exceeds $H(\vec{\psi})$ (and therefore DM must pay an associated large amount to their information provider). Intuitively, the more spread out the distribution of $\vec{\psi}$'s, the greater must be the mutual information in their channel for DM to achieve the same low value of expected regret (and therefore the more DM must pay their information provider to achieve that expected regret). On the other hand, if DM is willing to accept greater regret, $P(\epsilon = 1)$ and $H(\epsilon)$ will both be greater, and therefore Prop. 3 provides a smaller lower bound on $I(\vec{\psi}; \vec{t})$. So in this situation, DM need pay less to their information provider.

This can all be seen as a special case of Prop. 1. In particular, for $P(\epsilon = 1) \leq 1/k$, the lower bound $H(\vec{\psi}) - H(\epsilon) - P(\epsilon = 1)\ln_2[k]$ on $I(\vec{\psi}; \vec{t})$ given by Prop. 3(ii) is a convex and decreasing function of the regret ϵ . This agrees with the more generally applicable bound of Prop. 1.

In some scenarios, if \vec{t} is observable, then we can use Prop. 3 to make quantitative predictions concerning the behavior of DM. For example, we might be able to use historical data concerning pairs $(\vec{t}, \vec{\psi})$ to estimate the information-theoretic terms in Prop. 3. Assuming stationary probability distributions, we can then insert those estimates into Prop. 3 to compare the actual expected regret of DM, $\mathbb{E}(\epsilon) = P(\epsilon = 1)$, to the theoretically smallest possible regret. Conversely, by observing $\mathbb{E}(\epsilon)$, we can put bounds on the information-theoretic nature of the information channel that DM uses.

Example 3. Consider a DM whose job is to buy cloth for a major clothing manufacturer for the coming fashion season. This person must predict what color scheme will be in fashion for that coming season. For simplicity, we assume that only one such color scheme is in fashion for any single season. Let k be the number of possible color schemes. For simplicity, assume that all colors are 'equal', other than the arbitrary choice of nature as to which color will be in fashion, so that we can normalise payoffs such that the Savage utility to DM is 1 if their prediction of the color scheme matches what Nature actually does and 0 otherwise.

DM has some "observational" information concerning what the color scheme will be next fashion season, which they use to predict what element y will obey $\psi_y = 1$. Due to noise in that prediction, we model this as an information channel $P(\vec{t} | \vec{\psi})$ where \vec{t} is a k -dimensional binary vector exactly one of whose components equals 1. We interpret $P(\vec{t} | \vec{\psi})$ as the probability that DM predicts that $\text{argmax}_i(\{t_i\})$ will be color scheme next season when the actual color scheme of next

season is $\text{argmax}_i(\{\psi_i\})$. (Note that the number of possible values of \vec{t} is the same as the number of possible values of $\vec{\psi}$, namely k .)

If DM were able to use their observations to exactly predict the next color scheme, then $P(\vec{t} | \vec{\psi})$ would always equal 1 when evaluated for $\vec{t} = \vec{\psi}$. Since they cannot do this, $P(\vec{t} | \vec{\psi})$ in fact is spread out among multiple \vec{t} for any given $\vec{\psi}$. Intuitively, we would expect that the more “spread out” $P(\vec{t} | \vec{\psi})$ is, the worse the predicting of DM, and therefore the lower their expected utility. However as mentioned above, the prior $P(\vec{\psi})$ is also important. In fact, it is the relation between $P(\vec{\psi})$ and $P(\vec{t} | \vec{\psi})$ that we would expect to determine expected regret. For example, expected utility could be high even if $P(\vec{t} | \vec{\psi})$ were quite spread out for some $\vec{\psi}$, so long as it was only spread out for $\vec{\psi}$ that occur with low probability, while being tight for $\vec{\psi}$ that occur with high probability.

The formalization of this intuition is given by Prop. 3, in particular the first result there. Note in particular that for expected regret to be small, the LHS of Prop. 3(i) must be substantially less than $\ln[k]$. At the same time, $H(\vec{\psi} | \vec{t})$ is bounded above by $\ln[k]$. So for expected regret to be small, $H(\vec{\psi} | \vec{t})$ must be substantially less than its upper bound. Recall though that $H(\vec{\psi} | \vec{t})$ measures the accuracy of the information channel for those $\vec{\psi}$ that are most likely to occur; loosely speaking. So for expected regret to be small, the information channel must be quite accurate for those $\vec{\psi}$ that are most likely to occur.

3.4 Implications of the data-processing inequality

Return now to Ex. 1. Say that DM has already contracted with the polling company, so that the channel $P(\vec{t} | \vec{y}, \vec{\psi})$ is fixed. However say a second company now offers to sell a service to DM that will “clean” any vector of re-election probabilities the polling company might provide to DM, to “remove noise” and thereby make it more useful to DM. In particular, say that the second company claims that its procedure is guaranteed to map the \vec{t} 's DM receives from the polling company to new values of \vec{t} in a way that raises the correlation coefficient between \vec{t} 's and $\vec{\psi}$'s. DM decides to buy this service, presuming that if there is a higher correlation coefficient, then they will be better able to choose what bill to vote for. Can the second company be telling the truth? If so, can DM in fact make better choices of what bill to vote for by using this second company's services?

It turns out that the answer to the first question is 'yes'. It is possible to post-process the data coming out of an information channel to make it have higher correlation coefficient with the data going into the channel. However it turns out that it is *not* possible to post-process the data coming out of an information channel to make it have higher mutual information with the data going into the channel. As discussed below, this means that if DM is fully rational, and operating at the limits of rate distortion theory, the answer to the second question is 'no'.

To see that correlation coefficients between t and ψ can be raised by post-processing, it suffices to consider a simple example. Say that $k = 1$, and that the single value $\psi \in \{-2, -1, 0, 1, 2\}$. Assume that the prior probability over such values is $P(\psi) = [\delta_{\psi,1} + \delta_{\psi,-1}]/2$, where $\delta_{a,b}$ is the Kronecker delta function which equals 1 if $a = b$, and 0 otherwise. Presume that $t \in \{-2, -1, 0, 1, 2\}$, and that the information channel sets $t = \psi$ or $t = 2\psi$ with equal probability, independent of the value ψ . Presume also that the post-processing second company maps $t \rightarrow t$ if $|t| = 1$ and maps $t \rightarrow t/2$ otherwise. (In other words, they halve t if its magnitude is 2, and leave it alone if its

magnitude is 1.) Whether or not the second company intervenes, $\mathbb{E}(t) = \mathbb{E}(\psi) = 0$, and $\mathbb{E}(\psi^2) = 1$. Before the intervention of the second company, $\mathbb{E}(\psi t) = 3/4$, and $\mathbb{E}(t^2) = 5/2$. So before the intervention of the second company, the correlation coefficient equals $(3/4) \sqrt{2/5} \approx .47$. However after the intervention, the correlation coefficient equals 1. So the post-processing of the second company raises the correlation coefficient, exactly as they claimed.¹⁰ Notice though that the mutual information $I(\psi; t)$ is the same whether or not the second company intervenes. (It equals $\ln_2(2)$ in both cases.) So by the discussion above of rate distortion theory, there is no change in the minimal possible regret of DM.

This simple example illustrates a more general result. The **data-processing inequality** of information theory concerns any set of three random variables X, Y, Z taking values x, y, z respectively, where Z is statistically independent of X given Y , i.e., where $P(z | x, y) = P(z | y)P(y | x)$. It says that for any such triple, $I(Z; X) \leq I(Y; X)$ [5, 27]. In other words, say that a random variable Y is formed from a random variable X . Then no matter what kind of processing is done to $y \in Y$ to produce a new value $z \in Z$, it cannot result in Z 's having higher mutual information with X than does Y .

By the data-processing inequality, the value \vec{s} produced by the composite channel $\vec{y} \rightarrow \vec{t} \rightarrow \vec{s}$ produced by the second company's post-processing cannot have higher mutual information with $(\vec{\psi}, \vec{y})$ than does \vec{t} . If we write I_* for the maximal value of $I(\vec{S}; \vec{\Psi}, \vec{Y})$ over all possible kinds of such post-processing, then the data-processing inequality tells us that $I_* \leq I^*$. Therefore the set of composite channels (considered as conditional distributions stochastically mapping $(\vec{\Psi}, \vec{Y})$ into \vec{T}) with mutual information less than I_* is a subset of the set of channels with mutual information less than I^* .

Accordingly, if DM is fully rational, and always uses a hedonic utility U which results in the lowest possible expected regret given their information channel, they cannot decrease their expected regret by using that second company.

4 Restricted observation space \mathcal{T}

All of the information-theoretic results presented so far involve situations where \mathcal{T} is unrestricted, but P must obey certain constraints. An alternative restriction, considered by NR, is where \mathcal{T} is restricted, but there are no restrictions on P (other than the fact that it must map elements of Y to \mathcal{T}). For such problems the advantage of our framework is that it allows theoretical consideration of the assumptions made by NR, and illuminates the relationships with other decision-making models. For convenience, throughout this section we assume that the space of possible Savage utility values satisfies $\psi(Y) \subseteq [\underline{\psi}, \bar{\psi}]$, so that $\psi(Y)$ is bounded above and below.

We assume that DM can specify P and U subject to constraints on them. Therefore DM is faced with the problem of solving for the P and U that optimise the expected value of $\psi(y_i)$, where i is a randomly selected member of $\operatorname{argmax}_{i=1, \dots, k} U(t_i)$, and the expectation is taken over functions $\psi(\cdot)$ as well as $y_{1:k}$ and $t_{1:k}$, which are sampled according to the joint distribution defined by G_k and P .

¹⁰This phenomenon of post-processing the output of an information channel to improve its correlation with the input should not be confused with James-Stein shrinkage estimators, which also reduce quadratic error.

As an example of such a case, say that $P(\vec{t} | \vec{y}, \vec{\psi}) = \prod_{i=1}^k P(t_k | y_i)$ (the case in standard NR). Then if $k = 2$, the function to be optimised is

$$\int_{y_1 \in Y} \int_{y_2 \in Y} \int_{t_1 \in \mathcal{T}} \int_{t_2 \in \mathcal{T}} \left[G_2(dy_1, dy_2) P(dt_1 | y_1) P(dt_2 | y_2) \right. \\ \left. \times \{ \psi(y_1) \Theta(U(t_1) - U(t_2)) + \psi(y_2) \Theta(U(t_2) - U(t_1)) \} \right], \quad (6)$$

where $\Theta(x)$ is the Heaviside theta function taking values 1 if $x > 0$, 0 if $x < 0$ and $\frac{1}{2}$ if $x = 0$. If $k > 2$ it is more messy (and less instructive) to write down the explicit objective function — see Section 5. For the remainder of this section we focus on the case $k = 2$ considered by NR, and derive consequences of NR’s choice of restriction (finite \mathcal{T}) within the information-theoretic framework of this article; we see that several assumptions made in [32] are actually a consequence of optimizing expected Savage utility given the restriction.

In Appendix B we prove some basic results giving the form of the optimal observation channel P . In particular we prove in Appendix B.1 that a deterministic channel is optimal, so that there is no need to consider stochastic channels when performing the optimisations. This means that an RUM is not an optimal choice of $P(t | y)$ and U . Further restrictions on the channel, or a more complex optimisation problem, would be required for a rational designer of $P(t | y)$ and U to choose to use an RUM.

4.1 Comparing $k = 2$ options

We now consider exactly the case presented by NR, in which $k = 2$ options are independently and identically distributed from some distribution, and the decision space \mathcal{T} is finite. We show that the optimal channel in this case corresponds to a step function η which depends only on the sampled $\psi(y)$ values. This amounts to formally justifying the assumption made in NR, that hedonic utility is a deterministic step function of $\psi(Y)$ corresponding to a binning of $\psi(y)$ values. We then derive explicit results on the form of this step function, which provide further insights into the scenario presented by NR. The situation of a finite space $\psi(Y)$ of Savage utilities is also considered.

First note that the actual values of $U(t)$ in Eq. (6) are not important, merely the ordering of these values. Any strictly increasing function u can be applied to $U(t)$ and the set $\operatorname{argmax}_{i=1, \dots, k} U(t_i)$ will not change. Accordingly, the hedonic utility function in this framework is an ordinal quantity, i.e., it is a preference ordering, not a cardinal quantity. Hence, since \mathcal{T} is an arbitrary finite space, and the utility is only an ordinal quantity, without loss of generality we will assume that \mathcal{T} is the set of numbers $\{1, \dots, T\}$ and $U(t) = t$. We observe in Appendix B.1 that optimal decisions are made by choosing P to correspond to a single-valued map $\eta : Y \rightarrow \mathcal{T}$. We now derive the optimal form for η for the case where DM is presented with two options y_1 and y_2 that are generated by IID sampling. Define a “threshold map” to be a map η which divides $\psi(Y)$ into contiguous regions separated by decision boundaries:

$$\eta_{\tau}(y) \triangleq \min\{i | \tau_i \geq \psi(y)\}, \quad (7)$$

where $\tau = (\tau_0, \tau_1, \dots, \tau_T)$ satisfies $\underline{\psi} = \tau_0 \leq \tau_1 \leq \dots \leq \tau_T = \bar{\psi}$.

Lemma 4. Consider the $k = 2$ problem where $\psi(Y)$ is either finite or an interval of \mathbb{R} and \mathcal{T} is finite. Then a function of the form (7) is optimal.

Proof. The proof is technical and is given in Appendix B.2 □

Lemma 4 confirms NR's assumption that a threshold map is the appropriate type of map to consider. The only reason we cannot say that all optimal η functions satisfy this condition is because of complications involving sets of measure 0 in the case of $\psi(Y)$ being an interval.

This result, combined with the observations of Section 2, shows that the optimal channel corresponds to a step function η which depends only on the sampled $\psi(y)$ values. Hence we can concentrate on these sampled utility values, focussing on the distribution on $\psi(Y)$ induced by the independent sampling of options from Y according to G . These sampled utility values will be drawn according to some distribution (induced by G), and we write its density with respect to an appropriate measure μ on $\psi(Y)$ as f . (Here we assume μ is Lebesgue measure if $\psi(Y)$ is an interval and counting measure if it is discrete.)

In an abuse of notation, since it is a threshold map defined by Savage utility values, we shall also consider our channel function η to map the Savage utility space $\psi(Y)$ to \mathcal{T} in addition to mapping the option space Y to \mathcal{T} .

4.2 Decomposition of the loss function

Define \mathcal{L} to be the difference between the best possible expected utility using a noiseless information channel, and the actual expected utility that DM will receive using the noisy channel. The problem for DM is to choose τ to maximise the expected Savage utility they receive, which is equivalent to minimising \mathcal{L} . We begin our analysis in this subsection, by deriving an auxiliary result that shows how to decompose \mathcal{L} into expected losses between decision points. We will then use this result in subsequent subsections

Define the expected loss between decision points s and t as:

$$\mathcal{L}(s, t) = \int_{x_1=s}^t \int_{x_2=s}^t f(x_1)f(x_2) \left[\max\{x_1, x_2\} - \frac{1}{2}(x_1 + x_2) \right] \mu(dx_1)\mu(dx_2), \quad (8)$$

recalling that f is a density on the Savage utility space $\psi(Y)$ with respect to the measure μ .

Lemma 5.

$$\mathcal{L} = \sum_{i=1}^T \mathcal{L}(\tau_{i-1}, \tau_i). \quad (9)$$

Proof. With decision boundaries τ , the expected Savage utility is

$$\int \int f(x_1)f(x_2) \{ \Theta[\eta_\tau(x_1) - \eta_\tau(x_2)]x_1 + \Theta[\eta_\tau(x_2) - \eta_\tau(x_1)]x_2 \} \mu(dx_1)\mu(dx_2), \quad (10)$$

where $\Theta(z) = \begin{cases} 1 & \text{if } z > 0 \\ \frac{1}{2} & \text{if } z = 0 \\ 0 & \text{if } z < 0 \end{cases}$ is the Heaviside function. We can expand this integral as the sum of similar integrals restricted to intervals between the decision boundaries:

$$\begin{aligned} & \sum_{i=1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{i-1}}^{\tau_i} f(x_1)f(x_2)\frac{1}{2}(x_1+x_2)\mu(dx_1)\mu(dx_2) \\ & + 2 \sum_{i=1}^T \sum_{j=i+1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{j-1}}^{\tau_j} f(x_1)f(x_2)x_2\mu(dx_1)\mu(dx_2). \end{aligned} \quad (11)$$

(As a technical comment, if μ is counting measure, then the lower integration point in (11) is not included in the integral, whereas the upper integration point is, which is consistent with the definition of η .)

Also consider the expected Savage utility if perfect decision making were achieved:

$$\begin{aligned} & \int \int f(x_1)f(x_2)\max\{x_1, x_2\}\mu(dx_1)\mu(dx_2) \\ & = \sum_{i=1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{i-1}}^{\tau_i} f(x_1)f(x_2)\max\{x_1, x_2\}\mu(dx_1)\mu(dx_2) \\ & + 2 \sum_{i=1}^T \sum_{j=i+1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{j-1}}^{\tau_j} f(x_1)f(x_2)x_2\mu(dx_1)\mu(dx_2). \end{aligned} \quad (12)$$

The loss \mathcal{L} incurred as a result of choosing the decision boundaries τ is the difference between (12) and (11):

$$\begin{aligned} & \sum_{i=1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{i-1}}^{\tau_i} f(x_1)f(x_2)\left[\max\{x_1, x_2\} - \frac{1}{2}(x_1+x_2)\right]\mu(dx_1)\mu(dx_2) \\ & = \sum_{i=1}^T \mathcal{L}(\tau_{i-1}, \tau_i). \end{aligned} \quad (13)$$

□

4.3 A continuous space of Savage utility values

In this subsection we derive novel results concerning the case considered in NR, where the set of possible Savage utilities are an interval of \mathbb{R} , i.e. $\psi(Y) = [\underline{\psi}, \bar{\psi}]$, and f is a positive density function with respect to Lebesgue measure on this interval. We start by giving an explicit characterization of the relationship between consecutive decision boundaries under optimal binning:

Proposition 6. *Suppose the space $\psi(Y)$ is an interval of \mathbb{R} and f is a positive density function on $\psi(Y)$. Then*

$$\tau_i = \frac{\int_{\tau_{i-1}}^{\tau_{i+1}} xf(x)\mu(dx)}{\int_{\tau_{i-1}}^{\tau_{i+1}} f(x)\mu(dx)} = \mathbb{E}[x | x \in (\tau_{i-1}, \tau_{i+1})]. \quad (14)$$

Proof. Since $\psi(Y)$ is an interval, we are optimising \mathcal{L} over a compact space, and therefore the minimum is achieved. We can take partial derivatives of \mathcal{L} with respect to τ_i to find necessary conditions for the minimum. To calculate these partial derivatives, it helps to note that

$$\begin{aligned}\mathcal{L}(a, b) &= 2 \int_{x_1=a}^b \int_{x_2=a}^{x_1} f(x_1)f(x_2) \left[x_1 - \frac{1}{2}(x_1 + x_2) \right] \mu(dx_1)\mu(dx_2) \\ &= \int_{x_1=a}^b \int_{x_2=a}^{x_1} f(x_1)f(x_2)(x_1 - x_2) \mu(dx_1)\mu(dx_2),\end{aligned}$$

so that

$$\frac{\partial}{\partial b} \mathcal{L}(a, b) = \int_{x_2=a}^b f(b)f(x_2)(b - x_2) \mu(dx_2).$$

Similarly

$$\mathcal{L}(a, b) = \int_{x_1=a}^b \int_{x_2=x_1}^b f(x_1)f(x_2)(x_2 - x_1) \mu(dx_1)\mu(dx_2),$$

and

$$\frac{\partial}{\partial a} \mathcal{L}(a, b) = - \int_{x_2=a}^b f(a)f(x_2)(x_2 - a) \mu(dx_2)$$

We therefore see that

$$\begin{aligned}0 &= \frac{\partial \mathcal{L}}{\partial \tau_i} \\ &= \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + \frac{\partial \mathcal{L}(\tau_i, \tau_{i+1})}{\partial \tau_i} \\ &= \int_{\tau_{i-1}}^{\tau_i} f(\tau_i)f(x)(\tau_i - x) \mu(dx) - \int_{\tau_i}^{\tau_{i+1}} f(\tau_i)f(x)(x - \tau_i) \mu(dx) \\ &= \int_{\tau_{i-1}}^{\tau_{i+1}} f(\tau_i)f(x)(\tau_i - x) \mu(dx)\end{aligned}$$

Equation (14) follows directly. \square

In other words, τ_i is the the centre of mass of f restricted to the interval (τ_{i-1}, τ_{i+1}) . It is interesting to note that a similar rule is derived in [35]. There decision points are placed at evenly spaced percentiles of the utility distribution, and hence each decision point is the median of the distribution restricted to lying between the neighbouring decision points. Our result differs from this one due to our having a different objective function: in [35] the probability of error is minimised, whereas in this paper we minimise the expected loss due to errors.

4.4 S-shaped hedonic utilities

We can derive an S -shaped utility curve directly from Proposition 6, if we make the additional arbitrary assumption that utility is incremented by a fixed increment across each decision point τ_i . Given that assumption, whether the utility is S -shaped is directly related to the spacing of

the decision boundaries. Closer decision boundaries cause a steeper hedonic utility, and decision boundaries that are far apart cause a flatter hedonic utility. In light of this, consider the relationship (14). If f is increasing, then $\mathbb{E}[x|x \in (\tau_{i-1}, \tau_{i+1})] > \frac{1}{2}(\tau_{i-1} + \tau_{i+1})$, and so $\tau_{i+1} - \tau_i < \tau_i - \tau_{i-1}$ corresponding to the intervals shortening in length. This means that the hedonic utility is getting steeper in such a region. Similarly if f is decreasing then $\tau_{i+1} - \tau_i > \tau_i - \tau_{i-1}$ and the intervals are increasing in length. This means that the hedonic utility is getting flatter in such a region. Hence for a unimodal density f the optimal choices of P and U result in the hedonic utility being an S -shaped function of $\psi(y)$, with the bump in f corresponding to the middle part of the S .

Note though that this derivation of an S -shaped hedonic utility is entirely dependent on the choice of the precise utility values for $U(t)$, $t \in \mathcal{T}$. As noted above, we simply chose $\mathcal{T} = \{1, \dots, T\}$ and $U(t) = t$ for convenience. U could be chosen to be any increasing function of t without changing any of the optimality results presented, and hence the shape of the utility function is entirely arbitrary (subject to being monotonically increasing). Properly speaking, all this analysis (and Netzer's analogous analysis) shows is that the decision boundaries τ_i in $\psi(Y)$ are more closely spaced in regions of high density f . The standard way to derive cardinal utilities for a given social choice scenario is to extend that scenario, by introducing a lottery into it [39]. In Appendix C we show how this can be done for decision scenarios expressible in our framework. We defer detailed analysis of such decision scenarios involving lotteries for future work.

4.5 Finite $\psi(Y)$

In much of the experimental literature involving hedonic utilities, the set of options available to the decision maker is discrete (indeed finite). In particular, consider the scenario where Y is finite, having size M and resulting in the finite set $\psi(Y) \equiv \{x_1, \dots, x_M\}$. The standard calculations, including those carried out in the previous subsection, assume the existence of derivatives and therefore cannot be carried over to such a finite Y scenario. The framework we have described in Section 2 does not however limit us to considering the case of $\psi(Y)$ being a continuous interval of \mathbb{R} . In this subsection we show how one can use different techniques to solve for the optimal decision points for this scenario.

Recall that our problem is to choose the decision points to minimise the function \mathcal{L} given by (9). We formalize this problem as choosing $T - 1$ out of a set of M possible decision points, where each of those possible decision points lies half-way between the possible utility values x_i . So that set of possible decision points is $\{d_1, \dots, d_{M-1}\}$, where we define $\{d_i \triangleq \frac{x_i + x_{i+1}}{2} : i = 0, \dots, M\}$, with $d_0 = x_1 - 1$ and $d_M = x_M + 1$. The problem facing DM is how to choose $T - 1$ members of the set of possible decision points to be the decision boundaries τ_i . This problem can be solved as a shortest path problem in the following way.

Define the nodes of a directed graph to be (i, j) where $i = 1, \dots, T - 1$ corresponds to the decision boundaries τ_i , and $j = 1, \dots, M - 1$ correspond to possible decision boundaries d_j . An edge exists from node (i, j) to node (i', j') if $i' = i + 1$ and $j' > j$. In addition there is a node $(0, 0)$ with edges to each node $(1, j)$, and a node (T, M) which receives edges from each node $(T - 1, j)$. This network is shown in Fig. 1. The length of edge $(i, j) \rightarrow (i + 1, j')$ is $\mathcal{L}(d_j, d_{j'})$. A path from node $(0, 0)$ to (T, M) corresponds to a set of decision boundaries τ , where if the path goes through node (i, j) then the i th decision boundary τ_i is equal to the j th possible decision point d_j . The total

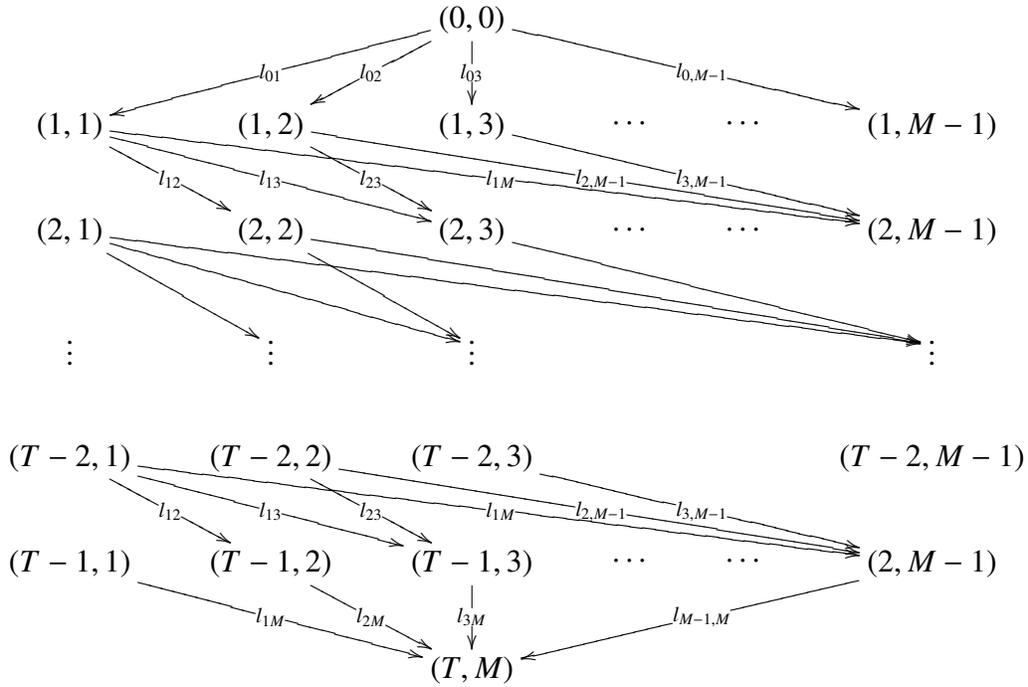


Figure 1: Dynamic programming network for finding optimal decision boundaries when $\psi(Y)$ is finite. The length of an edge from (i, j) to $(i + 1, j')$ is given by $l_{jj'} \triangleq \mathcal{L}(d_j, d_{j'})$. A path from $(0, 0)$ to (T, M) corresponds to a set of decision boundaries in $\psi(Y)$ — if the path goes through node (i, j) then the i th decision point τ_i is equal to the j th possible decision point d_j . The objective is to find the shortest path, which therefore minimises the total loss \mathcal{L} .

length of such a path is $\sum_{i=1}^T \mathcal{L}(\tau_{i-1}, \tau_i) = \mathcal{L}$. Hence the objective is to find a shortest path from $(0, 0)$ to (T, M) .

This problem is easily solved using dynamic programming. Define $s(i, j)$ to be the set of nodes which are at the end of an edge starting at (i, j) . We recursively set

$$V(i, j) = \min_{(i+1, j') \in s(i, j)} \left\{ \mathcal{L}(d_j, d_{j'}) + V(i+1, j') \right\}$$

with $V(i, j) = \infty$ if $s(i, j)$ is empty. The procedure will finish with an optimal \mathcal{L} value at node $(0, 0)$, and the optimal set of decision boundaries will correspond to the shortest path through the network. Note that the computational complexity of this procedure is $O(TM^2)$.

5 Comparing $k = 3$ observations

One might infer from the discussion in Section 4 that an implication of the model with restrictions on \mathcal{T} is the existence of an S -shaped hedonic utility function for any particular type of decision space Y (although one should also recall that the derived hedonic utility is actually only a preference ordering). However the analysis in that section follows NR in assuming that $k = 2$ options are compared. There is no reason why we should not also consider other values of k . After all, often in the real world a decision maker has more than 2 options. (Indeed, in RUM, $k = |Y|$, the extreme opposite case from NR.) To see what the implications of other choices of k might be, we can consider the case $k = 3$. We can then compare the results for this case to those obtained when $k = 2$, to see if the hedonic utility is affected by the number of presented alternatives.

As in Section 4, we assume that DM partitions $\psi(Y)$ into T regions, restrict attention to the case of $\psi(Y)$ being an interval of \mathbb{R} , and assume that the options presented are sampled independently. Thus the problem is identical to that considered in Section 4.3 but with $k = 3$ options presented instead of $k = 2$. As before, there is no loss in generality in assuming $\mathcal{T} = \{1, \dots, T\}$ and $U(t) = t$.

The analysis for this $k = 3$ case is more difficult than for the $k = 2$ case, since now DM can make mistakes in two ways. The first kind of mistake can arise when all three x_i 's result in the same t , i.e., fall in the same interval $[\tau_{i-1}, \tau_i]$. The second kind of mistake can arise when the two larger x_i 's fall in the same interval, while the remaining x_i falls in a lower interval.

Due to this extra complexity, we have not been able to find an optimal solution for the decision boundaries for the $k = 3$ scenario, as we could for the $k = 2$ scenario. However we can still compare the two scenarios. In particular, below we consider a special case, and show that the optimal decision boundaries for $k = 2$ and $k = 3$ must differ for that special case. Accordingly, the proper analysis would be to solve for the optimal decision boundaries under a probability distribution over k values, an analysis that is beyond the scope of this paper.

To simplify the calculations, we will perform all calculations conditional on the event $x_1 > x_2 > x_3$. By symmetry, the expected reward would be the same after conditioning on any other preference ordering, so there is no loss of generality in doing so. The conditional density of x_1, x_2 and x_3 is given by

$$f_{X_1, X_2, X_3 | X_1 > X_2 > X_3}(x_1, x_2, x_3) = \begin{cases} 6f(x_1)f(x_2)f(x_3) & \text{if } x_1 > x_2 > x_3 \\ 0 & \text{otherwise.} \end{cases}$$

The expected received utility given decision boundaries τ is therefore

$$\begin{aligned} \sum_{i=1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \left\{ \int_{x_2=\underline{\psi}}^{\tau_{i-1}} \int_{x_3=\underline{\psi}}^{x_2} 6f(x_1)f(x_2)f(x_3)x_1 \mu(dx_3) \mu(dx_2) \right. \\ \left. + \int_{x_2=\tau_{i-1}}^{x_1} \left[\int_{x_3=\underline{\psi}}^{\tau_{i-1}} 6f(x_1)f(x_2)f(x_3) \frac{x_1+x_2}{2} \mu(dx_3) \right. \right. \\ \left. \left. + \int_{x_3=\tau_{i-1}}^{x_2} 6f(x_1)f(x_2)f(x_3) \frac{x_1+x_2+x_3}{3} \mu(dx_3) \right] \mu(dx_2) \right\} \mu(dx_1). \end{aligned}$$

The first term in the square brackets corresponds to mistakes of the second kind, where x_3 is in a lower decision interval, whereas the second term in the square brackets corresponds to mistakes of the first kind, where all three x_i values are in the same decision interval.

A very similar decomposition of the expected utility under perfect decision making can be carried out, with the only difference being that both of the fractions are replaced by x_1 . Hence the loss due to imperfect decision making is given by

$$\begin{aligned} \mathcal{M} &= \sum_{i=1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{i-1}}^{x_1} \int_{x_3=\underline{\psi}}^{\tau_{i-1}} 6f(x_1)f(x_2)f(x_3) \left\{ x_1 - \frac{x_1+x_2}{2} \right\} \mu(dx_3) \mu(dx_2) \mu(dx_1) \\ &\quad + \sum_{i=1}^T \int_{x_1=\tau_{i-1}}^{\tau_i} \int_{x_2=\tau_{i-1}}^{x_1} \int_{x_3=\tau_{i-1}}^{x_2} 6f(x_1)f(x_2)f(x_3) \left\{ x_1 - \frac{x_1+x_2+x_3}{3} \right\} \mu(dx_3) \mu(dx_2) \mu(dx_1) \\ &= 3 \sum_{i=1}^T g(\underline{\psi}, \tau_{i-1}) \mathcal{L}(\tau_{i-1}, \tau_i) + 2 \sum_{i=1}^T \mathcal{M}(\tau_{i-1}, \tau_i), \end{aligned} \quad (15)$$

where $g(a, b) = \int_a^b f(x) \mu(dx)$ is the probability that x is between a and b , $\mathcal{L}(a, b)$ is as in equation (8), and

$$\begin{aligned} \mathcal{M}(a, b) &= \int_{x_1=a}^b \int_{x_2=a}^{x_1} \int_{x_3=a}^{x_2} f(x_1)f(x_2)f(x_3)(2x_1 - x_2 - x_3) \mu(dx_3) \mu(dx_2) \mu(dx_1) \\ &= \int_{x_3=a}^b \int_{x_2=x_3}^b \int_{x_1=x_2}^b f(x_1)f(x_2)f(x_3)(2x_1 - x_2 - x_3) \mu(dx_3) \mu(dx_2) \mu(dx_1). \end{aligned}$$

As before, we wish to choose τ to minimise \mathcal{M} . We differentiate with respect to τ_i , to get

$$\begin{aligned} \frac{\partial \mathcal{M}}{\partial \tau_i} &= 3 \left\{ g(\underline{\psi}, \tau_{i-1}) \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + f(\tau_i) \mathcal{L}(\tau_i, \tau_{i+1}) + g(\underline{\psi}, \tau_i) \frac{\partial \mathcal{L}(\tau_i, \tau_{i+1})}{\partial \tau_i} \right\} \\ &\quad + 2 \left\{ \frac{\partial \mathcal{M}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + \frac{\partial \mathcal{M}(\tau_i, \tau_{i+1})}{\partial \tau_i} \right\} \end{aligned}$$

As mentioned above, we have not been able to find an explicit solution as we did when $k = 2$. Instead we compare with a solution to the $k = 2$ case to see if a decision vector τ can be

simultaneously optimal for both $k = 2$ and $k = 3$. Let τ be a vector of decision boundaries that are optimal for $k = 2$. Recall from the proof of Proposition 6 that we therefore have

$$\frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + \frac{\partial \mathcal{L}(\tau_i, \tau_{i+1})}{\partial \tau_i} = 0$$

for all i . Note also that $g(\underline{\psi}, \tau_i) = g(\underline{\psi}, \tau_{i-1}) + g(\tau_{i-1}, \tau_i)$. Therefore

$$\begin{aligned} & g(\underline{\psi}, \tau_{i-1}) \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + g(\underline{\psi}, \tau_i) \frac{\partial \mathcal{L}(\tau_i, \tau_{i+1})}{\partial \tau_i} \\ &= g(\underline{\psi}, \tau_i) \left\{ \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + \frac{\partial \mathcal{L}(\tau_i, \tau_{i+1})}{\partial \tau_i} \right\} - g(\tau_{i-1}, \tau_i) \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} \\ &= -g(\tau_{i-1}, \tau_i) \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} \end{aligned}$$

and

$$\frac{\partial \mathcal{M}}{\partial \tau_i} = 3 \left\{ f(\tau_i) \mathcal{L}(\tau_i, \tau_{i+1}) - g(\tau_{i-1}, \tau_i) \frac{\partial \mathcal{L}(\tau_{i-1}, \tau_i)}{\partial \tau_i} \right\} + 2 \left\{ \frac{\partial \mathcal{M}(\tau_{i-1}, \tau_i)}{\partial \tau_i} + \frac{\partial \mathcal{M}(\tau_i, \tau_{i+1})}{\partial \tau_i} \right\}.$$

At this point, we consider the special case of $\psi(Y) = [0, 1]$ and $f(x) = 1$. Under this regime it is simple calculation to show that

$$\begin{aligned} g(a, b) &= b - a, \\ \mathcal{L}(a, b) &= \frac{(b - a)^3}{6}, \\ \mathcal{M}(a, b) &= \frac{(b - a)^4}{8}, \end{aligned}$$

and therefore

$$\begin{aligned} \frac{\partial \mathcal{M}}{\partial \tau_i} &= 3 \left\{ \frac{(\tau_{i+1} - \tau_i)^3}{6} - (\tau_i - \tau_{i-1}) \frac{(\tau_i - \tau_{i-1})^2}{2} \right\} + 2 \left\{ \frac{(\tau_i - \tau_{i-1})^3}{2} - \frac{(\tau_{i+1} - \tau_i)^3}{2} \right\} \\ &= -\frac{1}{2} \left\{ (\tau_i - \tau_{i-1})^3 + (\tau_{i+1} - \tau_i)^3 \right\} \end{aligned}$$

Since, from (14), $\tau_i = \frac{i}{T}$ if (as we are assuming) the τ_i are optimal for $k = 2$, we see that

$$\frac{\partial \mathcal{M}}{\partial \tau_i} = -T^{-3} \neq 0.$$

Hence the optimal τ for $k = 2$ is not optimal for $k = 3$, and in general the optimal decision boundaries depend on the number of options k to be compared.

In fact for this simple example it is basic algebra to show that

$$\frac{\partial \mathcal{M}}{\partial \tau_i} = 0 \quad \Rightarrow \quad \tau_i = \sqrt{\frac{\tau_{i-1}^2 + \tau_{i-1}\tau_{i+1} + \tau_{i+1}^2}{3}}. \quad (16)$$

This is in contrast with the solution for $k = 2$, in which

$$\tau_i = \frac{\tau_{i-1} + \tau_{i+1}}{2} \leq \sqrt{\frac{\tau_{i-1}^2 + \tau_{i-1}\tau_{i+1} + \tau_{i+1}^2}{3}}, \quad (17)$$

where the inequality is strict if $\tau_{i-1} \neq \tau_{i+1}$.

To illustrate the implications of this result, as in NR take $U(t) = t$ and in addition assume f is uniform. In this case, for $k = 2$, hedonic utility is a linear function of the Savage utility. However (16, 17) shows that for this case, if $k = 3$, and if in addition T is large, then the hedonic utility is a convex function of the Savage utility. This qualitatively different hedonic utility under the presentation of three options rather than two could be tested empirically using lotteries involving three options.

The fact that the optimal hedonic utility function varies with k implies that the relative values DM assigns to different options depends on the total number of options. On the other hand, say that DM does not know the value of k . In this case, the best DM can do is assume some distribution $P(k)$ of decision problems they encounter in their lifetimes, and solve for the (single) hedonic utility function that optimizes expected Savage utility under $P(k)$.

6 Conclusion

In this paper we have identified and started to analyse a general framework for decision making, which includes previous models as special cases. The new feature of our framework is that the restrictions which make the model interesting are inherent limitations associated with observational restrictions. We posit that observed decision-making should be optimal, subject to these constraints, and leave for the particular setting the question of how such optimality might occur. One route may be evolutionary optimality, although in our opinion the requirement of the decision-making problem to remain stationary over evolutionary timescales and the fact that selective pressure vanishes in the limits that give interesting results ($T \rightarrow \infty$) render this a particularly weak argument. Instead, an interpretation in which ψ is viewed as a conventional decision theory utility function (instead of as adaptive fitness) means that it is in DM's interest to maximise the expected value of ψ by definition. This could be achieved, subject to constraints on P and \mathcal{T} , through rational analysis or a learning procedure as well as by evolutionary means.

In our framework we do *not* define \mathcal{T} as values of a hedonic utility function. Rather they're the information theoretic way of quantifying observational limitations. Hedonic utility is then *assigned* by DM to the elements of \mathcal{T} , to provide DM a way to make choices. Accordingly, to be rational, DM must choose the hedonic utility that results in choices that are as good as possible, given the channel and G . In other words, in our framework hedonic utility is not an aspect of human behaviour given *ex ante*. Rather the very concept of hedonic utility is derived, as a way for DM to choose among elements of \mathcal{T} to optimise their choice.

By appealing to general results of information theory, we immediately observe some interesting economic consequences of the framework: we observe a fundamental rule of diminishing returns for additional investment in observational capability, a unique price that DM is prepared to pay for such observational capability, and a fundamental result that post-processing an information

channel cannot provide improved decision-making. In addition, in a particular example where there is a unique ‘correct’ choice we derived very general limitations on the maximal probability of a correct choice.

We then used our framework to derive some of the assumptions made by [32] as a result of optimality given a restricted observation space \mathcal{T} , and explored further some of the consequences. These include extensions to a discrete space of possible Savage utility values, which is directly relevant to the experimental literature, and to the situation where DM considers more than 2 options at once. However this is an extremely incomplete analysis of the possible implications of our framework. We outline three major issues that still need to be addressed.

We have observed that the entire framework can only predict hedonic preference orders, as opposed to cardinal hedonic utility functions. As in standard decision analysis, the obvious route to hedonic utility functions is by requiring DM to choose between lotteries. An initial re-casting of our framework into lotteries is given in Appendix C, where we show that when $|\mathcal{T}| = |Y|$, the optimal hedonic utility function equals the Savage utility function. However the more important case, as in the rest of the analysis of this paper, is when the channel places restrictions on the ability of DM (either by taking $|\mathcal{T}| < |Y|$ or through other channel restrictions). However we have not established whether the optimal hedonic utility function is unique in that case. Nor have we established whether it sometimes follows an *S*-shape, in accord with experimental data and the results of Netzer.

Another important issue is that we have either worked in general terms (in Section 3) or, in parallel with NR, we have only analysed an extremely crude model of limitations on the observational information channel: simply that the observation space \mathcal{T} is finite (in Sections 4 and 5). A more complete analysis would parallel that of the latter sections, but instead analyse optimality when the limitations on P are truly information-theoretic, such as restricting the maximal capacity of the information channel. This is left for future work.

A final issue that should be analysed in more detail is the difference in optimal decision boundaries observed when $k > 2$ (Section 5). In particular the case where one averages over k according to some distribution is consistent with the requirement to maximise expected reward. It may be that when confronted with a choice among $k > 2$ options, that humans always break their choice down to pair-wise comparisons. If this is the case, then analysis for $k = 2$ would suffice. However it is not at all clear that the experimental psychology literature establishes that humans only make pair-wise comparisons. (Indeed, RUM assumes humans compare all options in Y at once.)

Acknowledgements

The authors wish to thank Aram Harrow, University of Bristol, for help with formulating the dynamic programming solution for discrete payoff spaces. The research in this paper was supported by the ALADDIN project (Autonomous Learning Agents for Decentralised Data and Information Systems), which is jointly funded by a BAE Systems and EPSRC (Engineering and Physical Science Research Council) strategic partnership (EP/C548051/1).

A Collected notation

- Y — the space of options
- y, y', y_i — generic elements of Y
- N — the size of Y (if Y is finite)
- G — a distribution over Y
- k — the number of options y_i sampled IID according to G
- ψ — the Savage utility function $\psi : Y \rightarrow \mathbb{R}$
- $\psi(Y)$ — the set of possible Savage utility values
- $\underline{\psi}$ and $\bar{\psi}$ — the minimum and maximum possible ψ values (both assumed to be finite and well-defined)
- x — the sampled $\psi(y)$ value
- f — the density over $\psi(Y)$, with respect to a measure μ , that is induced by the distribution G over Y .¹¹
- V — Netzer's hedonic utility function $V : Y \rightarrow \mathbb{R}$.
- \mathcal{T} — the observation space of hedonic utility arguments
- t — a generic element of \mathcal{T}
- T — the size of \mathcal{T} if \mathcal{T} is finite
- U — the hedonic utility function $U : \mathcal{T} \rightarrow \mathbb{R}$.
- P — an observation channel specifying a conditional distribution on \mathcal{T} for each element of Y
- η — a map $Y \rightarrow \mathcal{T}$ corresponding to a deterministic channel
- $\tau = (\tau_0, \dots, \tau_T)$ — a vector of decision boundaries
- η_τ — the channel map corresponding to τ
- \mathcal{L} — the drop in expected utility when $k = 2$
- $\mathcal{L}(s, t)$ — the expect loss between s and t when $k = 2$
- \mathcal{M} and $\mathcal{M}(a, b)$ — correspond to \mathcal{L} and $\mathcal{L}(a, b)$ but for the case of $k = 3$

¹¹For example, if Y is a subinterval of the reals, $f(x) = \int dy G(y)\delta(\psi(y) - x)$.

- ϕ — a lottery (distribution) over Y
- $\tilde{\phi}$ — a lottery (distribution) over \mathcal{T}
- Θ — the Heaviside theta function

B Technical results

B.1 Deterministic channels are optimal

We wish to determine whether it suffices to consider a deterministic channel P , i.e., whether optimal decision-making might require a random channel. As we might expect, assuming (as we do here) that DM has complete freedom to set the channel, then DM should use a deterministic channel. (As mentioned, future work concerns analyzing the case where the information capacity of P is restricted.)

Lemma 7. *A deterministic channel P is optimal.*

A fully rigorous proof of this result requires consideration of measure theoretic issues beyond the scope this paper; we instead provide a sketch of a proof.

Proof. Recall that a decision instance of DM corresponds to sampling k IID copies of (y_i, t_i) where the marginal distribution of y_i is given by G and the distribution of t_i conditional on y_i is given by P . We can think of the sampling of $y_{1:k}$ and $t_{1:k}$ as a function from a sample space Ω to $(Y \times \mathcal{T})^k$. Assuming sufficient measurability, the sampling process is equivalent to considering a mapping $\eta : Y \rightarrow \mathcal{T}$ which is randomly selected, then sampling $y_{1:k}$, then setting $t_i = \eta(y_i)$. The “distribution” of η is given by the properties of the channel. Hence, defining R to be the received $\psi(y)$ value and $\hat{R} \triangleq \mathbb{E}(R)$,

$$\hat{R} = \mathbb{E}(R) = \mathbb{E}_{y_{1:k}}[\mathbb{E}[R | y_{1:k}]] = \mathbb{E}_{\eta}[\mathbb{E}_{y_{1:k}}[\mathbb{E}[R | y_{1:k}] | \eta]].$$

Assume P , and therefore the distribution over η 's, is optimal. If any deterministic η^* satisfies

$$\mathbb{E}_{y_{1:k}}[\mathbb{E}[R | y_{1:k}] | \eta = \eta^*] > \hat{R}$$

then P is not optimal. Hence, for any deterministic map $\tilde{\eta}$,

$$\mathbb{E}_{y_{1:k}}[\mathbb{E}[R | y_{1:k}] | \eta = \tilde{\eta}] \leq \hat{R}.$$

Therefore, for all $\tilde{\eta}$ in a set of measure 1 under the distribution induced by P ,

$$\mathbb{E}_{y_{1:k}}[\mathbb{E}[R | y_{1:k}] | \eta = \tilde{\eta}] = \hat{R},$$

since otherwise we would have $\mathbb{E}_{\eta}[\mathbb{E}_{y_{1:k}}[\mathbb{E}[R | y_{1:k}] | \eta]] < \hat{R}$. Taking \tilde{P} to be deterministic corresponding to any $\tilde{\eta}$ in this measure 1 set, we have constructed a deterministic optimal channel function. \square

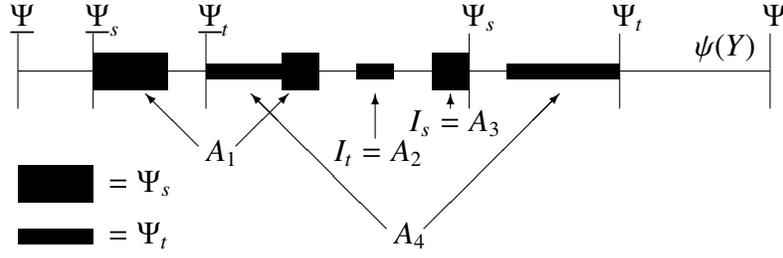


Figure 2: The subsets Ψ_s , Ψ_t , I_s , I_t , A_1 , A_2 , A_3 and A_4 used in the proof of Lemma 4

B.2 Step functions are optimal

Lemma 4. Consider the $k = 2$ problem where $\psi(Y)$ is either finite or an interval of \mathbb{R} and \mathcal{T} is finite. Then a function of the form (7) is optimal.

Proof. The definition (7) is equivalent to the statement that $\inf \psi(\eta^{-1}(t)) \geq \sup \psi(\eta^{-1}(t-1))$ for $t = 2, \dots, T$. We will prove that there exists an optimal η satisfying this latter condition.

In the case of uncountable $\psi(Y)$, subsets of $\psi(Y)$ of measure 0 under the sampling distribution of y will not affect the expected reward, and so we restrict attention to the case where $\psi(\eta^{-1}(t))$ is a union of intervals of positive measure under the sampling distribution. In the finite case $\psi(\eta^{-1}(t))$ is a finite union of contiguous subsets of $\psi(Y)$ with positive measure.

Suppose η is optimal, let $s = t - 1$, $\Psi_s = \psi(\eta^{-1}(s))$, $\underline{\Psi}_s = \inf \Psi_s$, $\bar{\Psi}_s = \sup \Psi_s$, and similarly for Ψ_t , $\underline{\Psi}_t$ and $\bar{\Psi}_t$. Note that since $\psi(Y) \subset [\underline{\psi}, \bar{\psi}]$ all of these quantities exist. They are shown in Figure 2.

We wish to prove that $\bar{\Psi}_s \leq \underline{\Psi}_t$. Suppose this is not the case, so that the interval $(\underline{\Psi}_t, \bar{\Psi}_s)$ is non-empty. Let I_s be the subinterval of Ψ_s with upper limit $\bar{\Psi}_s$, and let I_t be the subinterval of $\Psi_t \cap (\underline{\Psi}_t, \bar{\Psi}_s)$ closest to $\bar{\Psi}_s$. Again, these subsets are shown in Figure 2. By definition of I_t and I_s ,

$$\mathbb{E}(\psi(y) | \psi(y) \in I_t) < \mathbb{E}(\psi(y) | \psi(y) \in I_s).$$

Consider $\tilde{\eta}$ defined by

$$\tilde{\eta}(y) = \begin{cases} s & \text{if } \psi(y) \in I_t \\ \eta(y) & \text{otherwise.} \end{cases}$$

Then we can calculate $\mathbb{E}[R | \tilde{\eta}] - \mathbb{E}[R | \eta]$ by considering various locations of $\psi(y_1)$ and $\psi(y_2)$. Indeed the only time there is any difference in the expectations is when either $\psi(y_1) \in I_t$ and $\psi(y_2) \in (\Psi_s \cup \Psi_t) \cap I_t^c$ or vice versa. Hence write

$$A_1 = \Psi_s \cap I_s^c, \quad A_2 = I_t, \quad A_3 = I_s, \quad A_4 = \Psi_t \cap I_t^c$$

and let

$$p_i \triangleq \mathbb{P}(\psi(y_j) \in A_i), \quad m_i \triangleq \mathbb{E}(\psi(y_j) | \psi(y_j) \in A_i).$$

These sets are also shown in Figure 2.

Hence, by symmetry of y_1 and y_2 , and recalling that an option is selected at random if $\eta(y_1) = \eta(y_2)$,

$$\begin{aligned} & \mathbb{E}[R | \tilde{\eta}] - \mathbb{E}[R | \eta] \\ &= 2 \left\{ p_1 p_2 \left[\frac{m_1 + m_2}{2} - m_2 \right] + p_2 p_3 \left[\frac{m_2 + m_3}{2} - m_2 \right] + p_2 p_4 \left[m_4 - \frac{m_2 + m_4}{2} \right] \right\} \\ &= p_2 (p_1 m_1 + p_2 m_2 + p_3 m_3 + p_4 m_4 - (p_1 + p_2 + p_3 + p_4) m_2). \end{aligned}$$

Similarly, define $\tilde{\eta}$ by

$$\tilde{\eta}(y) = \begin{cases} t & \text{if } \psi(y) \in I_s \\ \eta(y) & \text{otherwise} \end{cases}$$

Through an identical argument, we see that

$$\mathbb{E}[R | \tilde{\eta}] - \mathbb{E}[R | \eta] = p_3 ((p_1 + p_2 + p_3 + p_4) m_3 - p_1 m_1 - p_2 m_2 - p_3 m_3 - p_4 m_4).$$

Since η is optimal, $\mathbb{E}[R | \tilde{\eta}] - \mathbb{E}[R | \eta] \leq 0$,

$$\begin{aligned} & \mathbb{E}[R | \tilde{\eta}] - \mathbb{E}[R | \eta] \leq 0 \\ \Rightarrow & p_1 m_1 + p_2 m_2 + p_3 m_3 + p_4 m_4 \leq (p_1 + p_2 + p_3 + p_4) m_2 \\ \Rightarrow & p_1 m_1 + p_2 m_2 + p_3 m_3 + p_4 m_4 < (p_1 + p_2 + p_3 + p_4) m_3 \\ & \text{since } m_2 < m_3 \text{ by construction} \\ \Rightarrow & \mathbb{E}[R | \tilde{\eta}] - \mathbb{E}[R | \eta] > 0, \end{aligned}$$

contradicting the assumption that η is optimal. □

C Hedonic utility instead of hedonic preference ordering

In Section 4.3 we observed that the choice of utility function U applied to the space \mathcal{T} is entirely arbitrary, and that the resulting optimal functional dependence of U on $\psi(y)$ is irrelevant up to monotonic transformations. To reinforce this point, note that if the difference in utility across decision boundary τ_i is $U(\tau_i^+) - U(\tau_i^-) = (\tau_{i+1} - \tau_i)^{-1}$, then U rises linearly with $\psi(y)$, rather than as the S shape analyzed in NR.

Stated differently, NR (and our analysis in this paper) only calculates an optimal hedonic preference, not an optimal hedonic utility. The way to get *cardinal* (not just ordinal) hedonic utility values would be to model a process of choosing between lotteries, giving a different model from NR.

As an illustration of how to model choosing between lotteries, let $\psi(Y)$ be a finite subset of $[0, 1]$, and ϕ_α be a probability distribution over Y , indexed by α . So $\{\phi_\alpha\}$ is a set of lotteries over Y . Without loss of generality, assume $\min_y \psi(y) = 0$ and $\max_y \psi(y) = 1$. The Savage expected utility associated with a lottery α is

$$\tilde{\psi}(\alpha) \triangleq \int \psi(y) \phi_\alpha(y) \mu(dy). \tag{18}$$

As before, we assume that DM cannot observe Y values directly, but due to observational limitations can only observe them through an information channel producing values in a finite space \mathcal{T} to which they assign hedonic utility values. As in the previous section and in NR, we will consider such channels that are equivalent to partitioning $\psi(Y)$ into T contiguous blocks, with the hedonic utility of block t increasing with the $\psi(Y)$ it contains.

Now any particular ϕ_α can be transformed from a lottery over Y it into a lottery over \mathcal{T} according to the rule

$$\begin{aligned}\tilde{\phi}(t|\alpha) &= \int P(t|y, \alpha)\phi_\alpha(y)\mu(dy) \\ &= \int P(t|y)\phi_\alpha(y)\mu(dy)\end{aligned}$$

under the assumption that t , the output of the information channel, only depends on the input to the channel, i.e. is conditionally independent of α given y . So if the hedonic utility function is written as $U(t)$ as before, then any particular α results in an expected hedonic utility of

$$\tilde{U}(\alpha) \triangleq \sum_t \tilde{\phi}(t|\alpha)U(t). \quad (19)$$

We can also rewrite (18) as an expectation over t values:

$$\tilde{\psi}(\alpha) = \sum_t \tilde{\phi}(t|\alpha)\mathbb{E}(\psi(y)|t, \alpha). \quad (20)$$

We assume there is some distribution of decision problems faced by DM, as before. However now that distribution is a distribution $P(\alpha)$ over the lotteries α , not over the consumption bundles $y \in Y$. DM will be given a sample of k lotteries $\{\alpha_1, \dots, \alpha_k\}$ formed by sampling $P(\alpha)$, and will choose the lottery $\text{argmax}_i[\tilde{U}(\alpha_i)]$. DM chooses both U and the decision boundaries defining $P(t|y)$ (which together set \tilde{U}) to result in as high a value of their Savage expected utility under this procedure as possible.

The distribution of lotteries on Y induces a distribution of lotteries over the finite set $\psi(Y)$, via (19). Here we assume that this distribution of lotteries over $\psi(Y)$ has full support on the associated simplex, so that every lottery over $\psi(Y)$ arises with non-zero probability.

Given this, if it is the case that

$$\tilde{U}(\alpha) > \tilde{U}(\beta) \quad \Leftrightarrow \quad \tilde{\psi}(\alpha) > \tilde{\psi}(\beta) \quad (21)$$

for all α and β then perfect decision making by DM will occur. Moreover if there exists α and β such that (21) is violated, then decision making is suboptimal, by the continuity of $\tilde{\psi}$ and \tilde{U} . Hence the decision making is optimal if and only if (21) holds. In general though, it is not possible to choose U so that (21) holds, due to the conditioning on α in the $\mathbb{E}(\psi(y)|t, \alpha)$ term of (20) and the restrictions imposed on the information channel P . So in general, decision-making is suboptimal.

However if $|\mathcal{T}| = |\psi(Y)|$ then each $\psi(y)$ value can be mapped to a unique $t \in \mathcal{T}$, so the conditioning on α disappears and the problem is to set hedonic utility values to each $\psi(y)$. As embodied in the following result, for this case (21) places restrictions on the cardinal values of $U(t)$. In contrast our earlier analysis and NR merely restrict the ordinal values of $U(t)$.

Proposition 8. *If $|\mathcal{T}| = |\psi(Y)|$, and the distribution over lotteries has full support on the set of all lotteries, then the optimal hedonic utility satisfies $U(t) = \psi(\eta^{-1}(t))$ for all $t \in \mathcal{T}$.*

Proof. Note from (19) that \tilde{U} is linear in U , so that non-negative scaling and translation of U does not change the relative ordering of any lotteries α and β . Therefore such scaling and translating will have no effect on the choices of DM. Hence we can assume, without loss of generality, that

$$\min_t U(t) = \min_y \psi(y) = 0 \quad \text{and} \quad \max_t U(t) = \max_y \psi(y) = 1.$$

We also assume, without loss of generality, that $\mathcal{T} = \{1, \dots, T\}$ and $\psi(\eta^{-1}(t))$ is an increasing function of t , so that the y 's giving least utility $\psi(y)$ are mapped to 1 and the y 's giving most utility $\psi(y)$ are mapped to T .

Suppose now that, for an arbitrary $t^* \notin \{1, T\}$, $U(t^*) > \psi(\eta^{-1}(t^*)) \in (0, 1)$. Consider a lottery α for which $\tilde{\phi}(T | \alpha) = 1 - \tilde{\phi}(1 | \alpha) = \psi(\eta^{-1}(t^*)) + \epsilon$. So

$$\begin{aligned} \tilde{\psi}(\alpha) &= (\psi(\eta^{-1}(t^*)) + \epsilon)\psi(\eta^{-1}(T)) + (1 - \psi(\eta^{-1}(t^*)) - \epsilon)\psi(\eta^{-1}(1)) = \psi(\eta^{-1}(t^*)) + \epsilon \\ \tilde{U}(\alpha) &= (\psi(\eta^{-1}(t^*)) + \epsilon)U(T) + (1 - \psi(\eta^{-1}(t^*)) - \epsilon)U(1) = \psi(\eta^{-1}(t^*)) + \epsilon \end{aligned}$$

since $\psi(\eta^{-1}(T)) = U(T) = 1$ and $\psi(\eta^{-1}(1)) = U(1) = 0$. Consider also a lottery β for which $\tilde{\phi}(t^* | \beta) = 1$. For this lottery

$$\begin{aligned} \tilde{\psi}(\beta) &= \psi(\eta^{-1}(t^*)) \\ \tilde{U}(\beta) &= U(t^*) > \psi(\eta^{-1}(t^*)) \end{aligned}$$

So, for sufficiently small ϵ , $\tilde{U}(\beta) > \tilde{U}(\alpha)$ and $\tilde{\psi}(\beta) < \tilde{\psi}(\alpha)$, and perfect decision-making is not achieved. A similar calculation can be carried out if $U(t^*) < \psi(\eta^{-1}(t^*))$.

On the other hand if we assume that $U(t) = \psi(\eta^{-1}(t))$ for all t then (21) holds and optimal decisions are made. \square

So we see that if individuals choose between lotteries instead of between individual options, we can go beyond the hedonic preferences framework to discover unique optimal hedonic utilities (modulo non-negative scaling and translation). In doing so, we discover that the S -shape of the hedonic utility as a function of Savage utility vanishes. It will be interesting to discover in future work if a unique optimal S -shaped utility function is recovered when we re-introduce the restriction $|\mathcal{T}| < |Y|$.

References

- [1] **Bester, H. and W. Güth.** 1998. "Is altruism evolutionarily stable?" *J. Econ. Behav. Organ.*, **34**: 193–209.
- [2] **Gintis, H., S. Bowles, R. Boyd, and E. Fehr,** ed. 2005. *The moral sentiments and material interests: The foundations of cooperation in economic life*. Cambridge, MA: MIT Press, 2005.

- [3] **Camerer, C. F.** 2003. *Behavioral game theory: experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- [4] **Camerer, C. F. and E. Fehr.** 2006. “When does “economic man” dominate social behavior?” *Science* **311**(5757): 47–52.
- [5] **Cover, T. M. and J. A. Thomas.** 1991. *Elements of information theory*. New York: Wiley.
- [6] **Dekel, E., J. Ely, and O. Yilankaya.** 2007. “Evolution of preferences.” *Rev. Econ. Stud.*, **74**(3): 685–704.
- [7] **Dreber, A., D. Rand, D. Fudenberg, and M. A. Nowak.** 2008. “Winners don’t punish.” *Nature*, **452**: 348–351.
- [8] **Fairhall, A. L., G. D. Lewen, W. Bialek and R. R. de Ruyter van Stevenick.** 2001. “Efficiency and ambiguity in an adaptive neural code.” *Nature* **412**: 787–792.
- [9] **Fehr, E. and U. Fischbacher.** 2003. “The nature of human altruism.” *Nature*, **425**: 785–791.
- [10] **Friedman, D.** 1989. “The S-shaped value function as a constrained optimum.” *Amer. Econ. Rev.*, **79**(5): 1243–1248.
- [11] **Fudenberg, D. and D. K. Levine.** 1998. *The theory of learning in games*. Cambridge, MA: MIT Press.
- [12] **Fudenberg, D. and J. Tirole.** 1991. *Game theory*. Cambridge, MA: MIT Press.
- [13] **Georgeff, M., B. Pell, M. Pollack, M. Tambe, and M. Wooldridge.** 1999. “The belief-desire-intention model of agency.” In *Intelligent Agents V: Agents Theories, Architectures, and Languages, Lecture Notes in Computer Science, vol. 1555.*, 1–10. Berlin/Heidelberg: Springer.
- [14] **Gurerk, O., B. Irlenbusch, and B. Rockenbach.** 2006. “The competitive advantage of sanctioning institutions.” *Science*, **312**(5770): 108–111.
- [15] **Güth, W.** 1995. “An evolutionary approach to explaining cooperative behavior by reciprocal incentives.” *Int. J. Game Theory*, **24**: 323–344.
- [16] **Hart, S.** 2005. “Adaptive heuristics.” *Econometrica*, **73**(5): 1401–1430.
- [17] **Hauert, C., A. Traulsen, H. Brandt, M. A. Nowak, and K. Sigmund.** 2007. “Via freedom to coercion: the emergence of costly punishment.” *Science*, **316**(5833): 1905–1907.
- [18] **Heifetz, A., C. Shannon, and Y. Spiegel.** 2007. “The dynamic evolution of preferences.” *Econ. Theory*, **32**(2): 251–286.
- [19] **Huck, S. and J. Oechssler.** 1999. “The indirect evolutionary approach to explaining fair allocations.” *Games Econ. Behav.*, **28**(1): 13–24.

- [20] **Wolpert, D.H. and J. Jamison.** 2011. “The Strategic Choice of Preferences: The Persona Model.” *Berkeley Electronic Journal of Theoretical Economics*.
- [21] **Wolpert, D.H. and J. Jamison.** 2011. “Schelling Formalized: Strategic choices of non-rational personas.” in Okasha, S. and K. Binmore (Ed’s), “Evolution and Rationality: Decisions, Cooperation and Strategic Behaviour”, *Cambridge University Press*.
- [22] **Kahneman, D.** 2003. “Maps of bounded rationality: Psychology for behavioral economics.” *Amer. Econ. Rev.*, **93**(5): 1449–1475.
- [23] **Keller, L. and H. K. Reeve.** 1998. “Familiarity breeds cooperation.” *Nature*, **394**: 121–122.
- [24] **Kreps, D. M., P. Milgrom, J. Roberts, and R. Wilson.** 1982. “Rational cooperation in the finitely repeated prisoners’ dilemma.” *J. Econ. Theory*, **27**(2): 245–252.
- [25] **Leland, J.** 1988. “A theory of approximate expected utility maximization.” Carnegie Mellon University Department of Social and Decision Sciences working paper.
- [26] **Luce, D.** 2003. “Whatever happened to information theory in psychology?” *Review of General Psychology*, **7**(2): 183–188.
- [27] **Mackay, D. J. C.** 2003. *Information theory, inference, and learning algorithms*. Cambridge, UK: Cambridge University Press.
- [28] **Manski, C. F.** 2005. “Random utility models with bounded ambiguity.” Paper presented at Berkeley summer symposium honoring Daniel McFadden, May 6-7, 2005. <http://elsa.berkeley.edu/symposia/econometrics05/papers/manski.pdf>.
- [29] **McFadden, D.** 1974. “Conditional logit analysis of qualitative choice behavior.” In *Frontiers in Econometrics*, ed. P. Zarembka, 105–142. New York: Academic Press.
- [30] **McNamara, J. M., Z. Barta, L. Fromhage, and A. I. Houston.** 2007. “The coevolution of choosiness and cooperation.” *Nature*, **451**: 189–192.
- [31] **Myerson, R. B.** 1991. *Game theory: Analysis of conflict*. Cambridge, MA: Harvard University Press.
- [32] **Netzer, N.** Forthcoming. “Evolution of time preferences and attitudes towards risk.” *Amer. Econ. Rev.*
- [33] **Nowak, M. A.** 2006. “Five rules for the evolution of cooperation.” *Science*, **314**(5805): 1560–1563.
- [34] **Rayo, L. and G. S. Becker.** 2007. “Evolutionary efficiency and happiness.” *J. Polit. Economy*, **115**(2): 302–337.
- [35] **Robson, A. J.** 2001. “The biological basis of economic behavior.” *J. Econ. Lit.*, **39**(1): 11–33.

- [36] **Rubinstein, A.** 1998. *Modeling bounded rationality*. Cambridge, MA: MIT press.
- [37] **Russell, S. J. and Subramanian, D.** 1995. “Provably bounded-optimal agents.” *Journal of AI Research*, **2**: 575–609.
- [38] **Samuelson, L.,** ed. 2001. Special issue on “Evolution of Preferences.” *J. Econ. Theory*, **97**(2).
- [39] **Savage, L. J.** (1954). *The Foundations of Statistics*. New York: John Wiley and Sons.
- [40] **Sharpee, T. and Bialek W.** 2007. “Neural decision boundaries for maximal information transmission.” *PLoS ONE* **2**(7): e646.
- [41] **Sims, C.** 2006. “Rational inattention: Beyond the linear-quadratic case.” *AEA Papers and Proceedings*: 158–163.
- [42] **Sims, C.** 2010. “Rational inattention and monetary economics.” In *Handbook of Monetary Economics*, ed. B. Friedman and F. H. Hahn, 155–181. Amsterdam: Elsevier.
- [43] **Starmer, C.** 2000. “Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk.” *J. Econ. Lit.*, **38**: 332–382.
- [44] **Thurstone, L. L.** 1927. “A law of comparative judgement.” *Psychological Review*, **34**(4): 278–286.
- [45] **Train, K. E.** 2003. *Discrete choice methods with simulation*. Cambridge, UK: Cambridge University Press.
- [46] **Trivers, R.** 2002. *Natural selection and social theory: Selected papers of Robert Trivers*. Oxford, UK: Oxford University Press.