

Mattsson, Lars-Göran; Weibull, Jörgen W.

Working Paper

Probabilistic Choice as a Result of Mistakes

IUI Working Paper, No. 544

Provided in Cooperation with:

Research Institute of Industrial Economics (IFN), Stockholm

Suggested Citation: Mattsson, Lars-Göran; Weibull, Jörgen W. (2000) : Probabilistic Choice as a Result of Mistakes, IUI Working Paper, No. 544, The Research Institute of Industrial Economics (IUI), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/95118>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

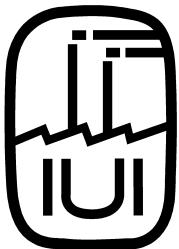
Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



THE RESEARCH INSTITUTE OF INDUSTRIAL ECONOMICS

Working Paper No. 544, 2000

Probabalistic Choice as a Result of Mistakes

by Lars-Göran Mattsson and Jörgen W. Weibull

IUI, The Research Institute of Industrial Economics
P.O. Box 5501
SE-114 85 Stockholm
Sweden

Probabilistic choice as a result of mistakes

LARS-GÖRAN MATTSSON* AND JÖRGEN W. WEIBULL
DEPARTMENT OF INFRASTRUCTURE AND PLANNING
ROYAL INSTITUTE OF TECHNOLOGY
SE-100 44 STOCKHOLM, SWEDEN
AND
DEPARTMENT OF ECONOMICS
STOCKHOLM SCHOOL OF ECONOMICS
P.O. Box 6501
SE-113 83 STOCKHOLM, SWEDEN

December 20, 2000

ABSTRACT. We derive a family of probabilistic choice models including the multinomial logit model, from a microeconomic model in which the decision maker has to make some effort in order to avoid mistakes when implementing any desired outcome. The disutility of this effort enters the decision maker's goal function in an additively separable way. A particular disutility function, yielding the multinomial logit and GEV models as special cases, is characterized axiomatically. Unlike the usual random-utility approach, the present approach leads to a normalization of the achieved utility with respect to the number of alternatives. The present model also applies to continuum choice sets in Euclidean spaces, and provides a microeconomic foundation for quantal response models in game theory.

JEL codes: D11, D60, D81

Keywords: choice, decision theory, mistakes

*The authors are grateful for helpful comments from PO Lindberg. Mattsson's research was given financial support from The Swedish Transport and Communications Research Board.

1. INTRODUCTION

In most real-life situations, the decision maker cannot guarantee any desired outcome with a probability exactly equal to one, only with a positive probability that may be close to one. The careless driver may drive off the road, the absent-minded shopper might buy the wrong item in the grocery store, the CEO who does not carefully monitor his staff may find that the firm has made suboptimal deals, etc. By contrast, in the classical microeconomic model, every economic agent can effortlessly obtain any desired choice alternative with certainty. In the present model, this is actually also possible for the decision maker. However, the marginal disutility of effort at zero mistake probability is assumed to be infinite, so the decision maker will optimally choose positive mistake probabilities.

Of course, there are probabilistic choice models in the economics literature, with the multinomial logit model as the prime example. These models are usually derived in an additive random-utility approach, with McFadden (1974) as the pioneering contribution. The randomness of choice is then interpreted either in terms of heterogeneity of tastes in a population of decision makers, or in terms of choice attributes hidden to the analyst, see Anderson *et al.* (1992). We instead consider a single decision maker who has to make some effort in order to avoid mistakes in the choice process. The higher probability he wants to put on any particular outcome, the more effort is needed. If he makes no effort, then the choice probabilities are given by an exogenous distribution, the “default” choice distribution. The decision maker is assumed to have deterministic preferences over the set of alternative outcomes, known by the analyst, and to be cognitively fully rational in the sense of being able to solve relevant maximization programs (or, at least, to behave as if he had this capacity). However, the model contains the above-mentioned element of procedural bounded rationality - there is a possibility of implementation mistakes.¹

Bounded procedural rationality of this type is not new in game theory. Selten (1975) defines “trembling hand” perfect equilibrium as a Nash equilibrium which is robust to small (exogenous) mistake probabilities when strategies are to be implemented (but players are cognitively fully rational). van Damme (1983) and van Damme and Weibull (1999) endogenize Selten’s mistake probabilities. This is also done here. However, we go beyond van Damme (1983) by providing an axiomatic characterization of a particular disutility function which generates a generalized logit choice model, and we complement the one-dimensional error-control model in van Damme and Weibull (1999) by providing an n -dimensional error-control model, where n is the number of discrete alternatives. The present model also provides a decision-theoretic founda-

¹Machina (1985) develops a model of stochastic choice at the individual level, based on deterministic (and known) preferences over lotteries. In this respect, our approach is similar to his.

tion for certain probabilistic choice behaviors in the “quantal response” approaches in Stahl (1990), Blume (1993, 1999), McKelvey and Palfrey (1995) and Chen *et al.* (1997). While these models start out by assuming probabilistic choice behavior, we provide conditions under which such behaviors are optimal for the decision maker. Moreover, while the latter models presume the parameters in the probabilistic choice formula for all players to be constant across the whole range of mixed-strategy profiles in the game in question - i.e. independent of the particular choice situation that is induced by the other players’ (expected) choice probabilities - we show when and why this constancy is consistent with expected utility maximization.

Technically, the discrete-choice version of the present model may be seen as a variant of the model in Proposition 3.7 in Anderson *et al.* (1992) (see also Fisk and Boyce (1984), Nadal *et al.* (1998) and Erlander (2000)). However, instead of the motivation in Anderson *et al.* of “variety-seeking behavior of the representative consumer” (p. 79), the present formulation builds on an explicit trade-off between expected payoff and disutility of mistake control, along with an axiomatic characterization of the disutility term. This results in a relatively general probabilistic choice model, which has as special cases the multinomial logit choice model, the generalized extreme-value (GEV) choice model, as well as the standard deterministic microeconomic utility maximization model of choice, and probabilistic generalizations of the latter. In the limit, as the weight given to the disutility of mistake control is taken to zero, the model boils down to the usual deterministic microeconomic choice model. Our formulation permits a resolution of the so-called “blue-bus-red-bus paradox,” and, unlike the usual random-utility approach to the logit model, the present approach leads to a normalization of the achieved utility with respect to changes in the number of alternative outcomes.² Except in special cases, the present choice model does not exhibit the independence of irrelevant alternatives property.

The rest of the paper is organized as follows. The basic model, including the axiomatic characterization of a particular control-cost function, is developed in section 2. The model is discussed and analyzed in section 3, and the multinomial and generalized extreme-value models are identified as special cases. Section 4 is devoted to an extension of the choice model to continuum sets of outcomes. Finally, some ideas for future research is discussed in section 5.

²Fisk and Boyce (1984) show how such a normalization can be achieved in the random-utility derivation of the logit model.

2. THE MODEL

First some notation. For any positive integer n , let Δ_n denote the $(n-1)$ -dimensional unit simplex in n -space,

$$\Delta_n = \left\{ p \in \mathbb{R}_+^n : \sum_{\omega=1}^n p_\omega = 1 \right\}, \quad (1)$$

and let $\text{int}(\Delta_n)$ be the relative interior of Δ_n in \mathbb{R}^n , $\text{int}(\Delta_n) = \Delta_n \cap \mathbb{R}_{++}^n$.³

We consider a decision maker who faces a decision problem with alternative outcomes $\omega \in \Omega = \{1, \dots, n\}$, where n is a positive integer, and where each outcome ω gives some *payoff* $\pi_\omega \in \mathbb{R}$. Let π be the associated vector of payoffs, $\pi = (\pi_1, \dots, \pi_n)$. We represent the decision maker's choice by a probability distribution $p = (p_1, \dots, p_n)$ over the outcomes. Hence $p \in \Delta_n$. The *expected payoff* under choice p is thus $p \cdot \pi = \sum_{\omega \in \Omega} p_\omega \pi_\omega$. If the decision maker without cost or effort can implement any choice $p \in \Delta_n$, then he would assign unit probability to the subset $\hat{\Omega} \subset \Omega$ of outcomes with maximal payoff, where

$$\hat{\Omega} = \{\omega \in \Omega : \pi_\omega = \hat{\pi}\} \quad (2)$$

and

$$\hat{\pi} = \max_{\omega \in \Omega} \pi_\omega. \quad (3)$$

The decision maker thus achieves the maximal payoff, just as in the standard micro-economic choice model.

Suppose, however, that there is a *disutility* or *control cost* $v(p, q)$ associated with every choice $p \in \Delta_n$, where $v : \Delta_n \times \text{int}(\Delta_n) \rightarrow \mathbb{R}_+$ is a continuous function satisfying the “no effort” condition

$$v(p, q) = 0 \quad \Leftrightarrow \quad p = q. \quad (4)$$

In other words, the decision maker has to make some effort when implementing a choice p , an effort that incurs a disutility or control cost, unless he chooses q , the *default choice vector* that results if the decision maker makes no effort.⁴

The (expected, total) *utility* associated with any choice $p \in \Delta_n$ is then defined as the expected payoff from the resulting outcome minus δ times the control cost, where δ is a positive scalar that represents the relative weight the decision maker attaches

³This terminology applies only to cases $n > 1$. In the case $n = 1$ it gives $\text{int}(\Delta_n) = \{1\}$.

⁴The notion of control costs in connection with strategy choice in games was introduced by van Damme (1983, 1991).

to the disutility or cost of implementation effort: $\pi \cdot p - \delta v(p, q)$. Hence, the new decision problem is

$$[M] \quad \max_{p \in \Delta_n} [\pi \cdot p - \delta v(p, q)] \quad (5)$$

The maximand is continuous in the choice variable p , and this variable is constrained to a compact set, so the decision problem has a non-empty compact solution set, by Weierstrass' Maximum Theorem. Given a payoff vector $\pi \in \mathbb{R}^n$ and disutility weight $\delta > 0$, let U^* denote the achieved utility level, and let P^* be the (non-empty and compact) solution set:

$$U^* = \max_{p \in \Delta_n} [\pi \cdot p - \delta v(p, q)] , \quad P^* = \arg \max_{p \in \Delta_n} [\pi \cdot p - \delta v(p, q)] . \quad (6)$$

By Berge's Maximum Theorem, this defines U^* as a continuous function of (π, δ) , and P^* as an upper hemi-continuous correspondence from pairs (π, δ) to (compact) subsets of Δ_n . Obviously, $\pi \cdot q \leq U^* \leq \hat{\pi}$. If all payoffs happen to be the same ($\pi_\omega = \pi_v$ for all $\omega, v \in \Omega$), then $U^* = \hat{\pi}$, and $P^* = \{q\}$ - it is optimal to make no effort. Otherwise, $U^* < \hat{\pi}$.

2.1. Axioms for disutility of mistake control. We now proceed to narrow down the class of disutility or control-cost function v , by way of imposing three axioms borrowed from information theory. For this purpose, we extend the domain of the cost functions to consist of all finite outcome sets Ω . Formally, now let the class of continuous control-cost functions be

$$v : \cup_{n \in \mathbb{N}} [\Delta_n \times \text{int}(\Delta_n)] \rightarrow \mathbb{R}_+ \quad (7)$$

that satisfy the “no effort” condition (4).

The first axiom requires that if control costs differ across outcomes, then that asymmetry should be fully captured by asymmetry in the default probability vector q . Hence, control cost should be invariant under any relabeling of the outcomes. By a *relabeling* of the outcomes in an outcome set Ω we mean a permutation of the set Ω , i.e., a bijection $\varphi : \Omega \rightarrow \Omega$. Formally:

[A1] (*Symmetry*) For any relabeling φ of Ω :

$$v [(p_1, \dots, p_n), (q_1, \dots, q_n)] = v [(p_{\varphi(1)}, \dots, p_{\varphi(n)}), (q_{\varphi(1)}, \dots, q_{\varphi(n)})]$$

The second axiom requires that the control cost associated with a choice that assigns equal probability to the m first outcomes, and zero to all others, in a situation

where the default distribution q is uniform over all outcomes, be decreasing in m and increasing in n , the number of outcomes in Ω . In other words, it is costly to completely avoid subsets of outcomes (to assign them zero probability), and more costly the larger the excluded subset is. This property can be formally expressed as follows. For any positive integers m and n such that $m \leq n$, let $r_n^m \in \Delta_n$ be the (“rectangular”) probability vector that assigns probability $1/m$ to each of the first m outcomes (and hence zero to all others).⁵

[A2] (*Monotonicity*) For any $m \leq n$, $v(r_n^m, r_n^n)$ is decreasing in m and increasing in n .

The third axiom requires that the cost of a choice p be independent of whether this is made directly or in two steps, by first choosing between two mutually exclusive and together exhaustive subsets of outcomes, and thereafter making a choice in the subset chosen in the first step. More exactly, by an *ordered binary decomposition* of the set Ω , granted this consists of two or more outcomes, we mean a pair (A, B) of non-empty and disjoint subsets of Ω , such that $A = \{1, \dots, m\}$ and $B = \{m + 1, \dots, n\}$. For any such decomposition, and any probability vector p in Δ_n , let $p_A = \sum_{\omega \in A} p_\omega$ and $p_B = \sum_{\omega \in B} p_\omega$ (with q_A and q_B defined similarly). The axiom states that the cost associated with choosing p directly, $v(p, q)$, equals the cost of the binary choice between A and B , $v[(p_A, p_B), (q_A, q_B)]$, plus the expected cost of choosing from the resulting subset A or B , where the first choice situation has probability p_A and the second probability p_B . In this second stage, the choice probabilities in the respective subset are the corresponding conditional probabilities.⁶

[A3] (*Decomposition*) For any decomposition (A, B) of Ω :

$$\begin{aligned} v(p, q) &= v[(p_A, p_B), (q_A, q_B)] + \\ &\quad + p_A v[(p_1/p_A, \dots, p_m/p_A), (q_1/q_A, \dots, q_m/q_A)] \\ &\quad + p_B v[(p_{m+1}/p_B, \dots, p_n/p_B), (q_{m+1}/q_B, \dots, q_n/q_B)] . \end{aligned}$$

It turns out that these axioms characterize the control-cost function, up to a positive scalar. This result is due to Hobson (1969).⁷

⁵By [A1], this monotonicity property holds irrespective of the order in which the outcomes are labelled.

⁶By [A1], this decomposition property holds for all binary decompositions of the outcome set.

⁷The result is a generalization of Shannon’s axiomatic characterization of the entropy function, see Shannon and Weaver (1949). For an alternative derivation of the relative-entropy function, see Snickars and Weibull (1977).

Proposition 1 [Hobson (1969)]. *Suppose v is a control-cost function. Then v satisfies [A1 – A3] if and only if, for some $\alpha > 0$,*

$$v(p, q) = \alpha \sum_{\omega \in \Omega} p_\omega \ln(p_\omega/q_\omega).$$

(Here and elsewhere we use the convention $0 \ln 0 = 0$.)

Let (A, B) be a binary decomposition of the outcome set Ω , as discussed in connection with axiom [A3]. Let the decision problem $[M']$ be to solve the maximization program (5) for the two aggregate outcomes A and B , and where the two payoffs are the achieved utilities when solving the maximization program (5) for $\Omega = A$ and $\Omega = B$, respectively. It can be verified that the two decision problems $[M]$ and $[M']$ are equivalent under axiom [A3], in the sense that they result in the same achieved utility and probability distribution over outcomes. In problem $[M']$, the probability of a particular outcome ω in the full set Ω is then calculated as the probability of the relevant subset A or B to which the outcome belongs, times the conditional probability of the outcome when chosen from within the relevant subset A or B .⁸ Conversely, if the cost function is such that for any binary decomposition of any decision problem $[M]$ the resulting problem $[M']$ is equivalent in this sense, then the control-cost function has to meet axiom [A3].

3. IMPLICATIONS

If the cost function meets axioms [A1] – [A3], then we may without loss of generality set

$$v(p, q) = \sum_{\omega} p_\omega \ln(p_\omega/q_\omega). \quad (8)$$

This is the relative-entropy (or information-gain) measure suggested by Kullback and Leibler, see Kullback (1959). This function is continuously differentiable in p on the relative interior of the simplex. A necessary condition for an interior solution p to the decision program $[M]$, for any such function v , is

$$\frac{\partial v(p, q)}{\partial p_\omega} = \frac{\pi_\omega + \lambda}{\delta} \quad \text{for all } \omega \in \Omega, \quad (9)$$

where λ is the Lagrangian multiplier associated with the constraint $\sum_{\omega} p_\omega = 1$. If the system of equations (9) has a solution $p \in \text{int}(\Delta_n)$, and $v(p, q)$ is convex in p , then such a solution is also a solution to the decision problem (5). In the present case (8), $v(p, q)$ is convex in p , and the first-order condition (9) becomes

⁸This equivalence, that follows from [A3], thus implies the choice axiom of Luce (1959).

$$\ln p_\omega + 1 - \ln q_\omega = \frac{\pi_\omega + \lambda}{\delta} \quad \text{for all } \omega \in \Omega, \quad (10)$$

with the unique solution p^* defined by

$$p_\omega^* = \frac{q_\omega \exp(\pi_\omega/\delta)}{\sum_{v \in \Omega} q_v \exp(\pi_v/\delta)} \text{ for all } \omega \in \Omega. \quad (11)$$

Since this solution is interior, it is the unique solution to the decision problem (5). In sum:

Proposition 2. *The unique solution to program [M] is $p^* \in \Delta_n$, defined in equation (11). Moreover, the achieved utility is*

$$U^* = \delta \ln \left[\sum_{\omega \in \Omega} q_\omega \exp(\pi_\omega/\delta) \right].$$

We note that if two outcomes have the same payoff, then their probability ratio equals the ratio of the two default probabilities. If one outcome has a lower payoff than another, then the probability ratio goes to zero when the weight $\delta > 0$ attached to the control-cost term is taken to zero. These two observations follow formally from the following equality:

$$\frac{p_\omega^*}{p_v^*} = \frac{q_\omega}{q_v} \exp\left(\frac{\pi_\omega - \pi_v}{\delta}\right). \quad (12)$$

In particular, in the limit when $\delta \rightarrow 0$ the solution p^* assigns unit probability to the subset $\hat{\Omega}$ of outcomes with maximal payoff, with probabilities in this subset proportional to the corresponding components of the default choice probability vector q . Hence, in the limit the decision-maker effortlessly achieves the maximal payoff $U^* = \hat{\pi}$, just as in the standard microeconomic choice model, and he does so by randomizing over the best outcomes according to the default choice probabilities.

We conclude with a comment on the present choice model (11) in relation to the so-called “blue-bus/red-bus paradox” in the discrete-choice literature (see, e.g., Anderson *et al.* (1992) or Ben-Akiva and Lerman (1985)). Suppose a decision maker faces the problem of whether to go by car, c , or by bus, b , for a certain trip. Let the payoffs be π_c and π_b , respectively. Let us call this binary choice decision problem 1. In decision problem 2, the bus alternative is replaced by two bus alternatives which differ only in some irrelevant aspect, such as the color of the bus (blue bus vs. red bus), and hence both outcomes, bb and rb , have the same payoff π_b . The paradox consists in the fact that the logit model (see next section) assigns lower probability

to the car outcome in decision problem 2 than in decision problem 1. However, in the present choice model this need not be the case. In particular, if the probability for the (aggregate) bus outcome in the no-effort case is unchanged ($q_{bb} + q_{rb} = q_b$), then the resulting choice probability for the car outcome is unchanged.

3.1. The logit model. Suppose all outcomes are equally likely in case the decision maker makes no effort, i.e., $q_\omega = \frac{1}{n}$ for all $\omega \in \Omega$. Then

$$v(p, q) = \ln n + \sum_{\omega} p_{\omega} \ln p_{\omega} . \quad (13)$$

Such a cost function, the negative of the entropy of the distribution p plus a constant, achieves its minimal value, zero, at the mid-point of the probability simplex, and its maximal value, $\ln n$, at each of its vertices (all probabilities but one equal to zero). By proposition 2 this results in the well-known (multinomial) *logit choice* model,

$$p_{\omega}^* = \frac{\exp(\pi_{\omega}/\delta)}{\sum_v \exp(\pi_v/\delta)} , \quad (14)$$

with achieved utility

$$U^* = \delta \ln \left[\frac{1}{n} \sum_{\omega} \exp(\pi_{\omega}/\delta) \right] . \quad (15)$$

In the random-utility foundation for the multinomial logit choice model, there is a random utility term u_{ω} associated with each outcome $\omega \in \Omega$, defined as the sum of a deterministic payoff π_{ω} (like here) and a random term ε_{ω} , where, by assumption, the random terms ε_{ω} are statistically independent and identically distributed according to the extreme-value distribution $F(x) = \exp[-\exp(-x/\delta)]$, for some $\delta > 0$. This leads to choice probabilities $p_{\omega}^* = \Pr(u_{\omega} \geq u_v \text{ for all } v \in \Omega)$ of precisely the form (14), see McFadden (1974) or Anderson et al (1992). In that framework, p_{ω} is thus the probability that ω is the outcome with the highest random utility $u_{\omega} = \pi_{\omega} + \varepsilon_{\omega}$ to the decision maker, and the parameter δ is proportional to the standard deviation of the random utility terms ε_{ω} .⁹ Hence, the weight attached to the disutility of control effort in the present model plays the same mathematical role as (the square root of) the variance of the random utility terms in the random-utility derivation of the multinomial logit model. A decision maker who more easily can control his or her choices correspond to a decision maker with less variance in the random utility terms.

Moreover, the present model gives a decision-theoretic justification of the practice in quantal-response models in game theory of letting this parameter be constant across

⁹More exactly, the standard deviation of ε_{ω} is $\delta\pi/\sqrt{6} \approx 1.28\delta$.

all strategy profiles, see e.g. McKelvey and Palfrey (1995). In such a game-theoretic setting, the expected payoffs π_ω to the player's own pure strategy ω depends on (the expectation of) the other players' mixed strategies.¹⁰ Given a player's preferences of the form suggested here, the induced optimal choice behavior is *as if* the random utility terms had the same variance irrespective of where in the polyhedron of mixed-strategy profiles the player expects the others to play.

How does the achieved utility in (15) depend on the number n of alternative outcomes? It was noted in section 2 that if all outcomes would happen to give exactly the same payoff, and another such outcome were added to the set Ω of alternatives, then the achieved utility would be unchanged: If $\pi_\omega = \hat{\pi}$ for all $\omega \in \Omega$, then $U^* = \hat{\pi}$, independently of n .¹¹ By contrast, in the above-mentioned random-utility derivation of the logit model (without control costs), the achieved utility would in this case be

$$E \left[\max_{\omega \in \Omega} (\pi_\omega + \varepsilon_\omega) \right] = \hat{\pi} + \delta(\gamma + \ln n) , \quad (16)$$

where γ is Euler's constant ($\gamma \approx 0.577$). In that derivation, an additional alternative increases the expected utility to the decision maker, even if the deterministic payoffs are identical, since the random term has a positive probability of being favorable to the decision maker. In the present model, however, the expected utility is unaffected. More generally, the addition of a new outcome in the present model results in higher (lower) achieved utility if and only if the payoff π_{n+1} to this additional outcome is higher (lower) than the originally achieved utility U^* . To see this, suppose that $\pi_{n+1} > U^*$, where $\pi = (\pi_1, \dots, \pi_n)$, and U^* is defined in equation (15). Then $\exp(\pi_{n+1}/\delta) > \exp[U^*/\delta]$, and thus

$$\frac{1}{n+1} \exp(\pi_{n+1}/\delta) > \frac{1}{(n+1)n} \sum_{\omega=1}^n \exp(\pi_\omega/\delta) = \left(\frac{1}{n} - \frac{1}{n+1} \right) \sum_{\omega=1}^n \exp(\pi_\omega/\delta) . \quad (17)$$

Re-arranging the terms, we obtain

$$\frac{1}{n+1} \sum_{\omega=1}^{n+1} \exp(\pi_\omega/\delta) > \frac{1}{n} \sum_{\omega=1}^n \exp(\pi_\omega/\delta) . \quad (18)$$

By taking the logarithm of both sides and multiplying by δ we see that addition of the new outcome indeed resulted in a higher achieved utility. (The same argument can

¹⁰Let ω be a pure strategy available to some player i , and let σ be a mixed-strategy profile in the game. The expected payoff to player i when playing ω against σ_{-i} is then a function of both ω and σ_{-i} .

¹¹This property holds for any control-cost function v that meets the “no effort” condition (4), see decision problem $[M]$.

be applied in the opposite case when $\pi_{n+1} < U^*$). Hence, if an additional alternative has high enough payoff, then this will more than compensate for the extra control cost.

3.2. The GEV model. By allowing for statistical dependence among the random terms in the random-utility derivation of the (multinomial) logit model, McFadden (1978, 1981) derived the probabilistic choice model,

$$p_\omega = \frac{e^{\pi_\omega} H_\omega(e^{\pi_1}, \dots, e^{\pi_n})}{H(e^{\pi_1}, \dots, e^{\pi_n})} \quad \forall \omega \in \Omega , \quad (19)$$

where $H : \mathbb{R}_{++}^n \rightarrow \mathbb{R}_{++}$ is a continuously differentiable and linearly homogeneous function with partial derivative H_ω with respect to its ω 'th argument. This is the so called *generalized extreme-value (GEV)* choice model. More precisely, he assumed that there is a random utility term u_ω associated with each outcome ω , defined as the sum of the deterministic payoff π_ω and a random term ε_ω , where the vector $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)$ of random terms has the cumulative probability distribution function $F(x_1, \dots, x_n) = \exp[-H(e^{-x_1}, \dots, e^{-x_n})]$.¹² This formulation allows for statistical dependence among the components of the random vector ε . McFadden (1978) showed that this leads to choice probabilities $p_\omega = \Pr(u_\omega \geq u_v \text{ for all } v \in \Omega)$ of the form (19). In this framework, p_ω is thus the probability that ω is the outcome with the highest random utility $u_\omega = \pi_\omega + \varepsilon_\omega$.

One can alternatively derive the choice model (19) in the present framework, as follows. For simplicity, let $\delta = 1$, and suppose the default choice vector q is a function of the payoff vector π , such that for all $\omega \in \Omega$

$$q_\omega = \frac{H_\omega(e^{\pi_1}, \dots, e^{\pi_n})}{\sum_v H_v(e^{\pi_1}, \dots, e^{\pi_n})} . \quad (20)$$

for some continuously differentiable and linearly homogeneous function $H : \mathbb{R}_{++}^n \rightarrow \mathbb{R}_{++}$. Using Euler's theorem, the optimal choice probabilities (11) can be written in the form (19), and the achieved utility becomes

$$U^* = \ln \left[\frac{H(e^{\pi_1}, \dots, e^{\pi_n})}{\sum_v H_v(e^{\pi_1}, \dots, e^{\pi_n})} \right] , \quad (21)$$

Note that this class of choice models includes the nested (multinomial) logit model, see, e.g., Ben-Akiva and Lerman (1985). We obtain the usual (multinomial) logit model in the special case when $H(x) = \sum_\omega x_\omega$.

¹²See Smith (1984) for necessary and sufficient conditions (in addition to H being homogeneous) for F to be a cumulative distribution function.

The above-mentioned “blue-bus/red-bus paradox” can be recast in terms of (20) for

$$H(x_c, x_{bb}, x_{rb}) = x_c + (x_{bb}^\sigma + x_{rb}^\sigma)^{1/\sigma}, \quad (22)$$

where $\sigma \geq 1$.¹³ In the boundary case $\sigma = 1$, all three outcomes are distinct, while the two bus alternatives become more and more indistinguishable as $\sigma \rightarrow +\infty$. With this particular function H (and recalling that the two bus outcomes have the same payoff) we obtain

$$q_c = \frac{1}{1 + 2^{1/\sigma}}, \quad q_{bb} = q_{rb} = \frac{2^{1/\sigma-1}}{1 + 2^{1/\sigma}} \quad (23)$$

For $\sigma = 1$ all three outcomes are equally likely, while $q_c \rightarrow 1/2$ and $q_{bb} = q_{rb} \rightarrow 1/4$ as $\sigma \rightarrow \infty$. In the special case when all payoffs are equal, we obtain by proposition 2, $p^* = q$. This indicates how the red-bus/blue-bus paradox may be resolved by parameterizing the extent to which the different outcomes are substitutes. Finally, we may note that when the payoffs for the car and bus outcomes differ, the achieved utility will be lower (higher) when the bus outcome is split into two differently colored outcomes, provided the bus payoff is lower (higher) than the car payoff. This difference will, however, tend to zero, independently of the levels of the payoffs, if the distinctness parameter σ tends to plus infinity.

4. CONTINUUM OUTCOME SETS

So far, we have assumed that the set Ω of outcomes is finite. However, part of the analysis can be generalized to a wide class of sets Ω , including such continuum sets in Euclidean spaces as the usual budget sets in the standard microeconomic consumer model. We here sketch one such extension.

More precisely, suppose that Ω is a non-empty compact set in some Euclidean space. Assume that the payoff function $\pi : \Omega \rightarrow \mathbb{R}$ is continuous. Let the decision maker’s choice be represented by a probability density function p on Ω . More exactly, let $p \in \Delta$, where Δ is the set of non-negative functions $p : \Omega \rightarrow \mathbb{R}$ with unit (Riemann) integral, $\int_{\Omega} p(\omega) d\omega = 1$. The expected payoff resulting from any choice $p \in \Delta$ is then $\int_{\Omega} \pi(\omega)p(\omega)d\omega$. Let $\text{int}(\Delta) \subset \Delta$ be the subset of functions $p \in \Delta$ that have a positive infimum on Ω . Just as in the previously analyzed case of finite Ω , we assume the existence of a “default choice” $q \in \text{int}(\Delta)$ and a “disutility” or “control cost” $v(p, q)$ associated with each choice $p \in \Delta$, where $v : \Delta \times \text{int}(\Delta) \rightarrow \mathbb{R}_+$. We restrict the subsequent analysis to the special case when v is the generalized Kullback-Leibler

¹³This inequality is needed in the random-utility formulation, but not in the present formulation.

measure¹⁴

$$v(p, q) = \int_{\Omega} p(\omega) \ln [p(\omega)/q(\omega)] d\omega . \quad (24)$$

The (expected total) utility associated with any decision $p \in \Delta$ is, as before, defined as the expected payoff minus δ times the disutility of mistake control. Hence, the current decision problem is

$$[M''] \quad \max_{p \in \Delta} \left[\int_{\Omega} \pi(\omega)p(\omega)d\omega - \delta \int_{\Omega} p(\omega) \ln [p(\omega)/q(\omega)] d\omega \right] . \quad (25)$$

In view of the solution (11) in the finite case, a natural solution candidate seems to be the function $p^* \in \text{int}(\Delta)$ defined by

$$p^*(\omega) = \frac{q(\omega) \exp [\pi(\omega)/\delta]}{\int_{\Omega} q(v) \exp [\pi(v)/\delta] dv} \quad \forall \omega \in \Omega . \quad (26)$$

This conjecture turns out to be true. (The following analysis also applies in the case when Ω is finite by letting the integration symbol represent summation.)

Proposition 3. *The function p^* defined by equation (26) solves program $[M'']$. Moreover, the achieved utility is*

$$U^* = \delta \ln \left(\int_{\Omega} q(\omega) \exp [\pi(\omega)/\delta] d\omega \right) .$$

Proof:¹⁵ Let \mathcal{F} be the set of (Riemann) integrable functions $f : \Omega \rightarrow \mathbb{R}_+$. Hence $\text{int}(\Delta) \subset \Delta \subset \mathcal{F}$. For each $\lambda \in \mathbb{R}$, let the functional $H_\lambda : \mathcal{F} \rightarrow \mathbb{R}$ be defined by

$$H_\lambda(f) = \int_{\Omega} (\pi(\omega) - \delta \ln [f(\omega)/q(\omega)] + \lambda) f(\omega) d\omega - \lambda . \quad (27)$$

One-variable calculus shows that the integrand is maximized at each ω if $f = f_\lambda^*$, where $f_\lambda^* \in \mathcal{F}$ is defined by $f_\lambda^*(\omega) = q(\omega) \exp ([\pi(\omega) + \lambda]/\delta - 1)$ for all $\omega \in \Omega$.¹⁶ Let $\mu \in \mathbb{R}$ be defined by

$$\exp (\mu/\delta - 1) = \left(\int_{\Omega} q(\omega) \exp [\pi(\omega)/\delta] d\omega \right)^{-1} . \quad (28)$$

¹⁴See Hobson (1969) for a discussion of the (24) continuum version of the Kullback-Leibler measure. Note that $v(p, q) \geq 0$, with equality if $p = q$. This follows from a point-wise application of the inequality $\ln(p/q) \geq 1 - q/p$ (an argument used in the finite case in Snickars and Weibull (1977)).

¹⁵The authors thank PO Lindberg for helpful suggestions.

¹⁶This follows from the fact that, for each $\omega \in \Omega$, the corresponding function $g : \mathbb{R}_+ \rightarrow \mathbb{R}$, defined by $g(x) = \pi(\omega)x - (\delta \ln [x/q(\omega)] + \lambda)x$, is continuously differentiable and concave with (unique) maximum at $x = q(\omega) \exp ([\pi(\omega) + \lambda]/\delta - 1)$.

Clearly $f_\mu^* \in \text{int}(\Delta) \subset \Delta$, and thus f_μ^* maximizes H_μ also over the subset $\Delta \subset \mathcal{F}$. But $H_\mu(f) = H_0(f)$ for all $f \in \Delta$, so f_μ^* also maximizes H_0 over Δ . This proves the claim, since $f_\mu^* = p^*$, and $H_0(f)$ is the maximand in $[M'']$. Finally, $U^* = H_0(p^*)$.

End of proof.

Similarly to the finite case, $\int_{\Omega} \pi(\omega)q(\omega)d\omega \leq U^* \leq \sup_{\omega \in \Omega} \pi(\omega)$. Just as in that case, the optimal choice probabilities concentrate on the set of outcomes with maximal payoffs when the disutility weight $\delta > 0$ is taken to zero, and the limiting choice probability density, inside the set of payoff-maximal outcomes, is proportional to the default probability density. Formally, if $\hat{\Omega} = \arg \max_{\omega \in \Omega} \pi(\omega)$, which is non-empty by assumption, then

$$p^*(\omega) \rightarrow \left[\int_{\hat{\Omega}} q(v)dv \right]^{-1} q(\omega)$$

for all $\omega \in \hat{\Omega}$ as $\delta \rightarrow 0$ (and thus $p^*(B) \rightarrow 1$ for any open set B containing $\hat{\Omega}$). Thus, in the limit we obtain the usual microeconomic model without implementation errors (modulo q).

We illustrate proposition 3 by way of an example. Let the outcome set be the unit interval, the payoff function linear, and the default choice the uniform density. Formally: $\Omega = [0, 1]$, with $\pi(\omega) = \omega$ and $q(\omega) = 1$ for all $\omega \in \Omega$. By proposition 3, the decision problem $[M'']$ is solved by

$$p^*(\omega) = \frac{1}{\delta[\exp[(1-\omega)/\delta] - \exp[-\omega/\delta]]} \quad (29)$$

and the achieved utility is

$$U^* = \delta \ln \delta + \delta \ln[\exp(1/\delta) - 1]. \quad (30)$$

As $\delta \rightarrow 0$, $p^*(\omega) \rightarrow 0$ for all $\omega < 1$, $p^*(1) \rightarrow \infty$ and $U^* \rightarrow 1$.¹⁷ This is expected: in the limit as control costs become negligable, the decision maker puts all choice probability at the payoff-maximizing outcome $\omega = 1$, and the obtained utility accordingly converges to the maximal payoff, $\pi(1) = 1$. See Figure 1 and 2 below. Note that $U^*(\delta) \rightarrow \frac{1}{2}$ when $\delta \rightarrow +\infty$. As the disutility of mistake control becomes overwhelming, there is no point in controlling mistakes, and the solution converges to the default choice q , which results in the expected payoff value $\frac{1}{2}$.

¹⁷To see that $p^*(\omega) \rightarrow 0$ for all $\omega < 1$, note that

$$\lim_{\delta \rightarrow 0} \delta e^{(1-\omega)/\delta} = \lim_{\theta \rightarrow +\infty} \frac{1}{\theta} e^{(1-\omega)\theta} = +\infty,$$

where the last equality follows from the fact that an exponential function eventually grows faster than a linear.

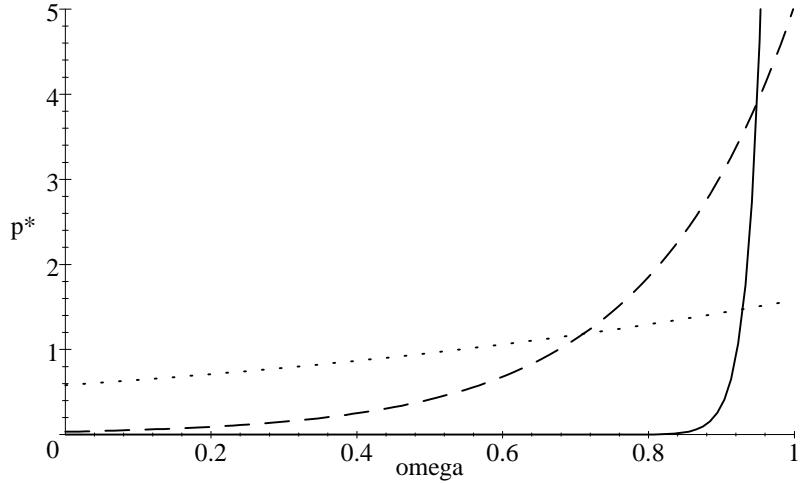


Figure 1. The optimal choice probability density function p^* , for $\delta = 0.02, 0.2$ and 1 .

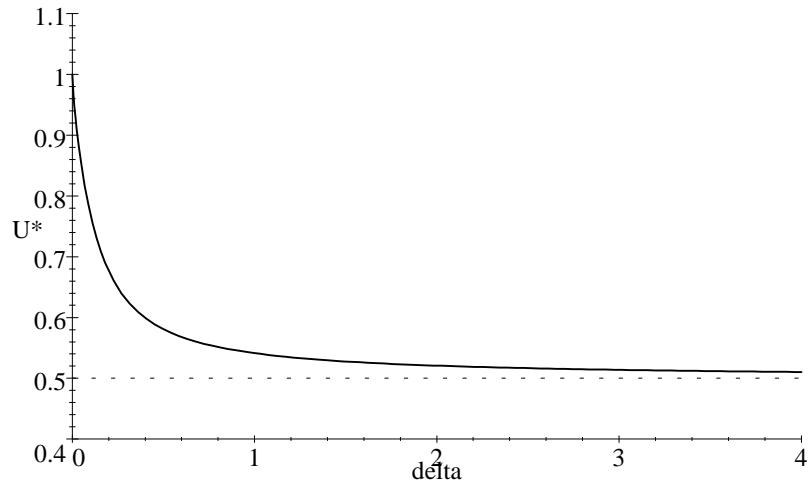


Figure 2: The achieved utility U^* as a function of δ .

5. CONCLUDING REMARKS

Our analysis has been focused on the case when the control-cost function is of the relative-entropy form (8). Then the optimal choice probabilities in decision problem $[M]$ are of the generalized logit form (11). A natural question then is whether other choice probability forms, e.g., the probit choice model, can be derived from some

other control-cost function. In the relative-entropy case, the control-cost function is convex in the choice probabilities. Hence, the achieved utility (in decision problem $[M]$) is the conjugate of a convex weighted control cost, implying that the achieved utility is convex in the pay-off vector π (see Rockafellar, 1970, p. 104).¹⁸ As has been noted by Hofbauer and Sandholm (2000) for the special case of a uniform default choice distribution q , this conjugate function value equals the expected utility in the random-utility derivation of a suitably normalized logit model. Moreover, the conjugate of this convex expected utility will retrieve the original weighted control-cost function.¹⁹ Hence, by taking the conjugate of the expected utility for other random utility models, candidate control-cost functions could possibly be derived.

Another avenue for further research might be to use the present framework to model social norms. Suppose that decision problem $[M]$ is faced by all individuals in a society in which defection from what others do causes disutility (see e.g. Lindbeck *et al.* (1999) for such a model). Letting the “default” decision q be the aggregate of all others’ decisions, and interpreting the disutility function v as the disutility of defecting from this aggregate behavior, decision problem $[M]$ determines each individual’s choice probabilities. A social equilibrium could then be thought of as a fixed point in this system of equations: an aggregate choice distribution q which results in individual choice probabilities (11) which, in aggregate, result in q .

REFERENCES

- [1] Anderson S., A. de Palma. and J.-F. Thisse (1992): *Discrete Choice Theory of Product Differentiation*. MIT Press (Cambridge, MA).
- [2] Baron R., J. Durieu, H. Haller and P. Solal (2000): “Control costs and potential functions for spatial games”, mimeo., CREUSET, University of Saint-Etienne.
- [3] Ben-Akiva M. and S. R. Lerman (1985): *Discrete Choice Theory: Theory and Application to Travel Demand*. MIT Press (Cambridge, MA).
- [4] Blume L. (1993): “The statistical mechanics of strategic interaction”, *Games and Economic Behavior* 5, 387-424.
- [5] Blume L. (1999): “How noise matters”, forthcoming, *Games and Economic Behavior*.

¹⁸In this particular context the domain and range of the cost function have been extended, by setting $v(p, q) = +\infty$ for all $p \in \mathbb{R}^n \cap (\sim \Delta_n)$, for any given $q \in \text{int}(\Delta_n)$.

¹⁹In the relative-entropy case with the convex control cost according to (8), the optimal value of $[M]$ is $U^*(\pi) = \delta \ln \sum_{\omega \in \Omega} q_\omega \exp(\pi_\omega / \delta)$ (see proposition 2). It is easily shown that the conjugate of this convex function is the weighted control cost, i.e., $\max_{\pi} [\pi \cdot p - U^*(\pi)] = \delta \sum_{\omega \in \Omega} p_\omega \ln(p_\omega / q_\omega)$.

- [6] Chen H.-C., J. Friedman and J.-F. Thisse (1997): “Boundedly rational Nash equilibrium: A probabilistic choice approach”, *Games and Economic Behavior* 18, 32-54.
- [7] van Damme E. (1983): *Refinements of the Nash Equilibrium Concept*. Lecture Notes in Economics and Mathematical Systems 219, Springer-Verlag (Berlin).
- [8] van Damme E. (1991): *Stability and Perfection of Nash Equilibria*. 2nd edition. Springer-Verlag (Berlin).
- [9] van Damme E. and J. W. Weibull (1999): “Evolution and refinement with endogenous mistake probabilities”, The Research Institute of Industrial Economics, Stockholm, WP No 525.
- [10] Erlander S. (2000): “Welfare, freedom of choice and composite utility in the logit model”, mimeo., Department of Mathematics, Linköping University.
- [11] Fisk C. S. and D. E. Boyce (1984): “A modified composite cost measure for probabilistic choice modeling”, *Environment and Planning A* 16, 241-248.
- [12] Hobson A. (1969): “A new theorem of information theory”, *Journal of Statistical Physics* 1, 383-391.
- [13] Hofbauer J. and W. Sandholm (2000): “Evolution and learning in games with randomly disturbed payoffs”, mimeo., Department of Economics, University of Wisconsin.
- [14] Kullback S. (1959): *Information Theory and Statistics*. Wiley (New York).
- [15] Lindbeck A., S. Nyberg and J. Weibull (1999): “Social norms and economic incentives in the welfare state”, *Quarterly Journal of Economics* 64, 1-35.
- [16] Luce R. D. (1959): *Individual Choice Behavior: A Theoretical Analysis*. Wiley (New York).
- [17] Machina M. (1985): “Stochastic choice functions generated from deterministic preferences over lotteries”, *Economic Journal* 95, 575-594.
- [18] McFadden D. (1974): “Conditional logit analysis of qualitative choice behavior”, in P. Zarembka (ed.), *Frontiers in Econometrics*. Academic Press (New York), 105-142.

- [19] McFadden D. (1978): “Modelling the choice of residential location”, In A. Karlqvist *et al.* (eds.) *Spatial Interaction Theory and Planning Models*. North Holland (Amsterdam), 75-96.
- [20] McFadden D. (1981): “Econometric models of probabilistic choice”, in C. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data with Econometric Applications*. Harvard University Press (Cambridge, MA), 198-272.
- [21] McKelvey R. and T. Palfrey (1995): “Quantal response equilibria for normal form games”, *Games and Economic Behavior* 10, 6-38.
- [22] Nadal J.-P., O. Chenevez, G. Weisbuch and A. Kirman (1998): “A formal approach to market organization: choice functions, mean field approximation and maximum entropy principle”, mimeo., GREQAM, Université d’Aix-Marseille.
- [23] Rockafellar R. T. (1970): *Convex Analysis*. Princeton University Press (Princeton, NJ).
- [24] Selten R. (1975): “Reexamination of the perfectness concept for equilibrium points in extensive games”, *International Journal of Game Theory* 4, 25-55.
- [25] Shannon C. and W. Weaver (1949): *The Mathematical Theory of Communication*. University of Illinois Press (Urbana, IL).
- [26] Smith T. E. (1984): “A choice probability characterization of generalized extreme value models”, *Applied Mathematics and Computation* 14, 35-62.
- [27] Snickars F. and J. W. Weibull (1977): “A minimum information principle”, *Regional Science and Urban Economics* 7, 137-168.
- [28] Stahl D. (1990): “Entropy control costs and entropic equilibria”, *International Journal of Game Theory* 19, 129-138.