Research in Economics 70 (2016) 197–231

Contents lists available at ScienceDirect

Research in Economics

journal homepage: www.elsevier.com/locate/rie

How to model consumer heterogeneity? Lessons from three case studies on SP and RP data

Michael P. Keane [a,*], Nada Wasi [b]

[a] Department of Economics, University of Oxford, UK
[b] Survey Research Center, Institute for Social Research, University of Michigan, USA

ARTICLE INFO

Article history:
Received 2 February 2016
Accepted 25 February 2016
Available online 4 March 2016

Keywords:
Discrete choice models
Consumer demand
Consumer heterogeneity
Mixture models
Generalized multinomial logit
Discrete choice experiments

ABSTRACT

The structure of consumer taste heterogeneity in discrete choice demand models is important, as it drives the structure of own and cross-price elasticities of demand, and the pattern of competition between products. Here we compare performance of three leading discrete choice models, using three datasets with very different properties. The models are the mixed logit with normal heterogeneity (N-MIXL), the generalized multinomial logit (G-MNL) and the mixture-of-normals logit (MM-MNL). Which model is preferred depends on the context: G-MNL does an excellent job of capturing the sort of departures from normality that are prevalent in stated preference (SP) data. But MM-MNL can capture more general departures from normality that are prevalent in revealed preference (RP) data. The finding that the structure of consumer taste heterogeneity is very different in SP vs. RP data suggests that caution should be applied before using SP to answer questions about the distribution of taste heterogeneity in actual markets. In an application to RP data on demand for frozen pizza, we obtain the interesting result that when a variety of a brand raises its price, most of the lost market share goes to other brands (rather than alternative varieties of the same brand). This suggests modeling heterogeneity in tastes for varieties is quite important for understanding brand switching.

1. Introduction

The question of how best to model consumer preference heterogeneity in discrete choice demand analysis has been a subject of great interest (and controversy) for at least 40 years. The reason the issue is so important is that the taste heterogeneity distribution drives the pattern of competition between products, as well as the structure of own and cross-price elasticities of demand. These features of a market matter for two main reasons:

(1) In deciding whether to introduce a new product, a firm must predict both the market share of the product and where it will come from (e.g., will the new product steal customers away from other firms, or just from the firm's own existing products?). The same question arises with respect to the change in market share if a firm changes the price of an existing product;

* Correspondence to: University of Oxford, Nuffield College, New Road, Oxford, OX1 1NF, UK.
E-mail address: michael.keane@nuffield.ox.ac.uk (M.P. Keane).

http://dx.doi.org/10.1016/j.rie.2016.02.002
1090-9443/

(2) In deciding whether to allow a merger of two firms, an antitrust regulator is concerned with the degree of competition between their products. If their products have high cross-price elasticities of demand, a merger may generate large price increases in equilibrium.

Unfortunately, the discrete choice model that has traditionally been the most widely used, the simple multinomial logit (MNL) with a common coefficient vector for all consumers (i.e., no heterogeneity), cannot address issues like (1) and (2). This is because MNL imposes *a priori* that all market shares move proportionately when a new product is introduced in a market (or when an existing product changes its price). This is obviously unrealistic.

For example, if Toyota introduces a new small economy car, we would expect it to compete with other small economy cars, and have little impact on family cars or luxury cars. But MNL makes the unrealistic prediction that the new car will steal market share proportionally from all alternatives. This problem can be solved by developing models where consumers are *heterogeneous* in their tastes for product attributes (such as fuel economy, seating capacity, performance, etc.). This allows cross-price elasticities to be determined flexibly by the data.

Since about 2000, it would be fair to say that the most popular discrete choice demand model among sophisticated practitioners and academics has been the so-called mixed logit model (MIXL). This is an extension of MNL that allows for population heterogeneity in the parameter vector. In most applications it has been assumed to be normally distributed in the population, leading to the logit model with normal mixing, or "N-MIXL" model. Early papers that developed and advocated this approach were Berry (1994), Berry et al. (1995), Revelt and Train (1998), Harris and Keane (1999) and McFadden and Train (2000). N-MIXL is popular because it allows for much more flexible patterns of substitution among alternatives than MNL, yet it is quite simple to estimate using simulation methods, and it is available in standard packages.

Another branch of the literature advocates the use of discrete mixtures-of-normals distributions to model consumer heterogeneity. In this model, the population is assumed to consist of a discrete set of types, each with its own normal heterogeneity distribution. Keane and Wasi (2013) refer to a mixed logit model where the heterogeneity distribution is specified as a discrete mixture-of-normals as the "mixed-mixed-logit" or MM-MNL model.

The virtue of the discrete mixtures-of-normals approach is both theoretical and practical. Theoretically, as shown by Ferguson (1973), a discrete mixture-of-normals can approximate any heterogeneity distribution arbitrarily closely. Practically, as shown in Geweke and Keane (1997, 2001) and Rossi et al. (2005), just a small number of mixture components (i.e., 2 or 3) can usually approximate even very complex heterogeneity distributions quite well.

Despite these appealing features, the discrete mixture-of-normals approach to modeling heterogeneity has not achieved anything close to the popularity of the normal-mixture-of-logits (N-MIXL) approach. We suspect this is because discrete mixture-of-normals models are more complex to estimate, tend to proliferate parameters, and require judgement in choosing the number of mixture elements. For these reasons, for better or worse, they have remained largely the preserve of advanced econometricians.

In a recent contribution, Fiebig et al. (2010) proposed a new model called the generalized multinomial logit or G-MNL model. This is a mixed logit model where the heterogeneity distribution is specified as a scaled mixture-of-normals. That is, the multivariate normal coefficient vector of the N-MIXL model is scaled by a positive scalar random variable, which Fiebig et al. (2010) assume is log normally distributed in the population. Of course, in the logit model, a scaling of the entire coefficient vector is observationally equivalent to an inverse scaling the idiosyncratic errors. Thus, one interpretation of the G-MNL model is that the idiosyncratic errors (or taste shocks) are more important (relative to the observed attributes) for some consumers than others.[1] This is often called "scale heterogeneity."

The Fiebig et al. (2010) paper was influential, and the G-MNL model became very popular among practitioners almost as soon as it was introduced.[2] There are a number of reasons for this sudden popularity of G-MNL: First, the model nests N-MIXL, reducing to that model as the variance of the scale heterogeneity parameter goes to zero. Second, the model adds only two parameters to the N-MIXL model (as we discuss below), and it is hardly any more difficult to estimate.[3] Third, Fiebig et al. (2010) showed in several applications that the G-MNL model usually gave a much better fit to consumer choice behavior than the N-MIXL model. If a model is no more difficult to estimate than its main competitor, is strictly more general, and often provides a substantially better fit, it is hard to see why it would not be popular.

Keane and Wasi (2013) provided further support for the G-MNL model by comparing it to the theoretically more general discrete mixture-of-normals approach on ten example data sets. We found that in most instances G-MNL is preferred over MM-MNL by standard metrics of model fit that penalize proliferation of parameters. This is because the G-MNL model typically gives almost as good a fit as MM-MNL, but it does so with many fewer parameters.

The point of the present paper is to provide some guidance as to the typical context in which each of these three models – N-MIXL, G-MNL or MM-MNL – is the preferred approach to modeling heterogeneity. We present three case studies, and assess which model performs best in each case. Much more importantly, we explain *why* each model emerges as preferred

[1] Fiebig et al. (2010) interpret their model as incorporating both "taste" and "scale" heterogeneity, although some may prefer to interpret the scaled mixture-of-normals as simply another parametric heterogeneity distribution, and not adopt this behavioral interpretation.

[2] For example, Shugan (2014) ranked Fiebig et al. (2010) as the 3rd most cited marketing article of the current decade, and the most cited methodological article. See *Shugan's Top 20 Marketing Meta-Journal*, April 2014, at: http://bear.warrington.ufl.edu/centers/MKS/.

[3] This contrasts with discrete mixture-of-normal models, where proliferation of parameters and choice of the number of elements of the mixture are both serious issues for the practitioner.