ITLS

WORKING PAPER

ITLS-WP-08-04

Revealing the extent of process heterogeneity in choice analysis: An empirical assessment

By

Sean Puckett & David Hensher

January 2008

This research is funded by the Australian Research Council Discovery Project Grant DP0770618

ISSN 1832-570X

INSTITUTE of TRANSPORT and LOGISTICS STUDIES

The Australian Key Centre in
Transport and Logistics Management

The University of Sydney

Established under the Australian Research Council's Key Centre Program.

NUMBER:	Working Paper ITLS-WP-08-04
TITLE:	**Revealing the extent of process heterogeneity in choice analysis: An empirical assessment**

ABSTRACT:

Choice analysts increasingly use a mix of revealed preference and stated choice data paradigms to identify preferences of samples of individuals that are used to infer behavioural response and willingness to pay for specific attributes. These data are in a sense artificial constructs that are developed to approximate real choice settings of the way that individuals process relevant information in making choices. As such all data designs formalized through a survey instrument seek information through questions that become *descriptions* of events and as such the probabilities of choice that are of interest are strictly probabilities attached to event descriptions and not choice probabilities of events per se. The recognition of this distinction, initially noted by Kahneman *et al.* (1982), can be captured, at least in part, through the idea of process heterogeneity as a way of recognizing and accounting for the many ways in which individuals process information, and in part is influenced by the way the analyst describes the context in which preference data is sought. Building on previous contributions on attribute processing, this paper draws on recent empirical evidence to further reinforce the importance of joint modelling of process and outcome in choice analysis. This study adds to the evidence of a trend emerging on the upward bias of mean estimates of marginal willingness to pay when ignoring process heterogeneity.

KEY WORDS: ***Process heterogeneity, event description, stated choice, design complexity, relevance, subadditivity, packaging***

AUTHORS: Sean Puckett & David Hensher

CONTACT: Institute of Transport and Logistics Studies **(C37)**

The Australian Key Centre in Transport Management

The University of Sydney NSW 2006 Australia

Telephone:	+61 9351 0071
Facsimile:	+61 9351 0088
E-mail:	itlsinfo@itls.usyd.edu.au
Internet:	http://www.itls.usyd.edu.au

DATE: January 2008

1. Introduction

Choice modelling has advanced significantly in recent years in terms of the econometric representation of preference heterogeneity within a sample (see Train 2003) and in terms of ways in which data are obtained, notably the design of choice experiments (see Rose and Bliemer 2007) and the tailored mechanisms available to capture behaviourally rich choice responses through internet and computer aided personal survey instruments (CAPI) (see Hensher *et al.* 2007).

What has been given less attention, although by no means ignored (see Bonini *et al.* 2004 and Hensher 2008), has been the questioning of the appropriateness of the frameworks used to obtain preference information from a sample of respondents (Rabin 1998). Specifically, the nature of the stated or revealed choice setting may not be sufficiently 'realistic' to provide the necessary information required to obtain choice probabilities and estimates of willingness to pay for specific attributes in the context of other attributes (and their levels). The recognition of process heterogeneity (Hensher 2008) suggests that more effort should be invested in understanding the relationship between what is offered up to each respondent and how they process such information, including the extent to which the experimental or revealed setting (neither of which is a true definition of an actual event) increases the 'gap' between the choice probability of interest and the 'choice' probability that is obtained. This 'gap' or bias is in part a reflection of a potential over-simplification of the way that the heuristics that individuals adopt in choice making in real markets are clouded by the use of *descriptions* of events instead of events per se.

The ability to portray events in the exact way that individuals perceive real world events is virtually impossible in a deterministic sense, and the best the analyst can do is to try and represent the events as realistically as possible, including the possibility of event variation over time (e.g., through new products and/or levels of attributes of existing products). What the analyst is unable to do adequately through a focus on choice outcome is to identify the influence that specific process strategies have on outcome. Specifically, individuals bring a whole raft of processing (including judgmental) capability and rules to the table when faced with scenarios designed with great statistical ingenuity by analysts, and it is the judgments of evidence strength or support that underlies numerical judgments of preference (be it through a first preference, a rank or a rate) and subsequent influence on choice probability.

Building on the contribution by Hensher and colleagues on attribute processing and the important distinction between complexity and relevancy, this paper draws on recent empirical evidence to further reinforce the importance of joint modelling of process and outcome in choice analysis. This research considers the effects of processing heterogeneity utilised by respondents for *every alternative in every choice set* faced, acknowledging that varying process rules may be enacted not only across decision makers, but across choice tasks faced by a given decision maker.

2. Empirical context to assess the presence of process heterogeneity

The data utilised in the empirical discussion within this section are from a 2005 study of road freight transport providers and their customers in Sydney, Australia. The study was designed to elicit preferences under a hypothetical road user charging system. The first step in the process involved administering the experiment to representatives of freight transport firms. Centred on a CAPI survey with a d-optimal experimental design (discussed in Puckett and Hensher 2007), the stated choice experiment involves three distinct procedures: (1) non-stated-choice questions intended to capture the relevant deliberation attributes and other contextual effects; (2) choice menus corresponding to a freight-contract-based setting (see Figure 1 and Puckett and Hensher 2007); and (3) questions regarding the attribute processing strategies enacted by respondents within each choice set. The resulting estimation sample, after controlling for outliers and problematic respondent data[1], includes 108 transporters, yielding 432 choice sets. The response rate was 45%.

In all cases except for the variable charges, the attribute levels for each of the stated choice alternatives are pivoted off of the levels of the reference alternative, as detailed below. The levels are expressed as deviations from the reference level, which is the exact value specified in the corresponding non-stated choice questions, unless noted:

Free-flow time: -50%, -25%, 0, +25%, +50%

Congested time: -50%, -25%, 0, +25%, +50%

Waiting time at destination: -50%, -25%, 0, +25%, +50%

Probability of on-time arrival: -50%, -25%, 0, +25%, +50%, with the resulting value rounded to the nearest five percent (e.g., a reference value of 75% reduced by 50% would yield a raw figure of 37.5%, which would be rounded to 40%). If the resulting value is 100%, the value is expressed as 99%. If the reference level is greater than 92%, the pivot base is set to 92%. If the pivot base is greater than 66 percent (i.e., if 1.5 times the base would be greater than 100%) let the pivot base equal X, and let the difference between 99% and X equal Y. The range of attribute levels for on-time arrival when X > 66% are (in percentage terms): X-Y, X-.5*Y, X, X+.5*Y, X+Y. This yields five equally-spaced attribute levels between X-Y and 99%.

Fuel cost: -50%, -25%, 0, +25%, +50% (representing changes in fuel taxes of -100%, -50%, 0, +50%, +100%)

(6) Distance-based charges: Pivot base equals .5*(reference fuel cost), to reflect the amount of fuel taxes paid in the reference alternative. Variations around the pivot base are: -50%, -25%, 0, +25%, +50%

[1] Preliminary analysis revealed that the degree of heterogeneity in reference trips was sufficiently high that some outliers obscured the inferential power of the data. After careful consideration, the following observations were removed from the final sample: (a) trips based on a fuel efficiency over 101 litres per 100 kilometres (or approximately twice the average fuel consumption for the larger trucks in the sample); (b) trips based on a probability of on-time arrival less than 33 percent; (c) round trips (or tours) of less than 50 kilometres; and (d) round trips of more than 600 kilometres. The trips eliminated, based on low fuel efficiency, may have obscured the results due to significantly prohibitive values for fuel cost and variable charges, reflecting reference trips that are too atypical to be pooled with other trips. An alternative source of obscuring effects via low fuel efficiency may be that the implied values of fuel efficiency were inaccurate, and hence either made the trade-offs implausible to respondents or reflect an inability of the respondent to offer meaningful information on which to base the alternatives. The trips eliminated, based on low probability of on-time arrival, are likely to have obscured the results because the trips involved travel quality significantly worse than the remainder of the sample, making the pooling of these trips into the sample problematic. Similarly, extremely short or long trips may have involved trade-offs that are significantly different to the trade-offs made by respondents in the sample at large.

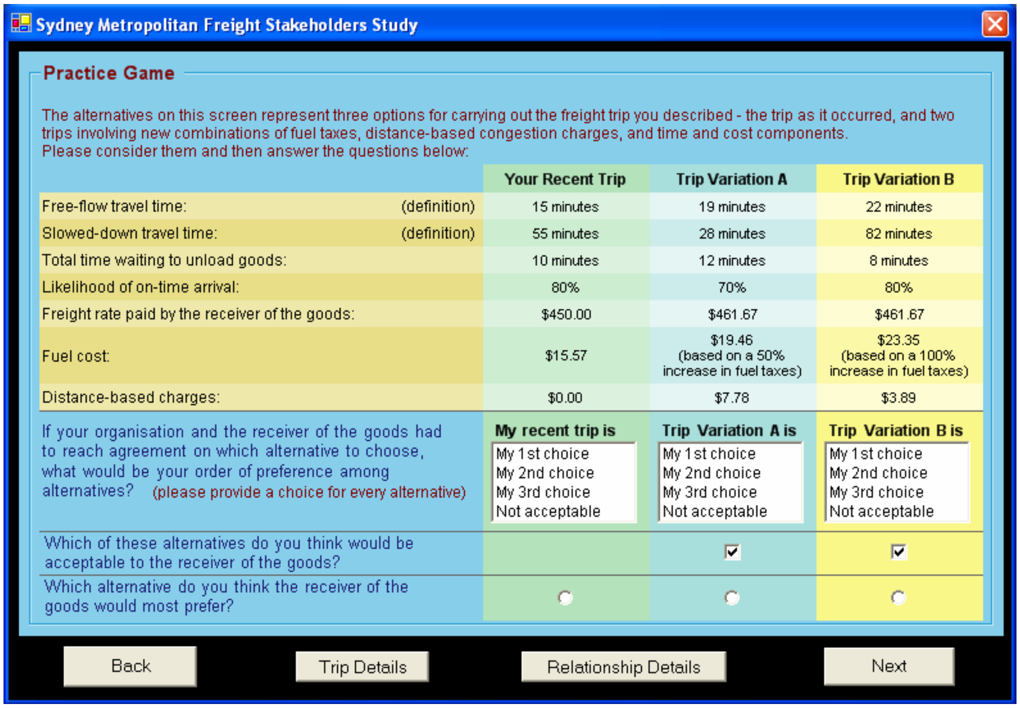

Figure 1: Main choice set screen

The extant literature, with rare exception (see Hensher *et al.* 2006, Hensher 2008) either ignores process heterogeneity altogether, utilises the information structure within choice sets as an indicator of cognitive burden, or uses global (i.e., across all choice sets faced) indicators of attribute exclusion and agglomeration. This research considers the effects of attribute processing strategies utilised by respondents for *every alternative in every choice set* faced, acknowledging that varying process rules may be enacted not only across decision makers, but across choice tasks faced by a given decision maker. The data set is to our knowledge unique in this sense.

To condition our empirical model (see Hensher *et al.* 2007) on attribute processing information, the stated attribute processing strategies (APSs) of each respondent *for each attribute within each alternative* faced were utilised in a transformation of the data. Following the independent choices made in each choice set, respondents are asked to reveal components of the APS utilised when making their choices. An explanation is offered by prefacing the first APS specification task with the statement (Figure 2), "Is any of the information shown not relevant when you make your choice? If an attribute did not matter to your decision, please click on the label of the attribute below. If any particular attributes for a given alternative did not matter to your decision, please click on the specific attribute. You may click on a selected item to de-select it."

This transformation involved four related behavioural stages: (1) the assignment of marginal utilities of zero to all attributes that were indicated as ignored; (2) the assignment of marginal utilities of zero to all *individual attributes* that were indicated as aggregated; (3) the inclusion of aggregate measures of transit time (i.e., the sum of free-flow time, slowed-down time and waiting time) and the transporter's costs (i.e., the sum of variable charges and fuel cost); and (4) the assignment of marginal utilities of zero to any aggregate measures that were not formed by the respondent.

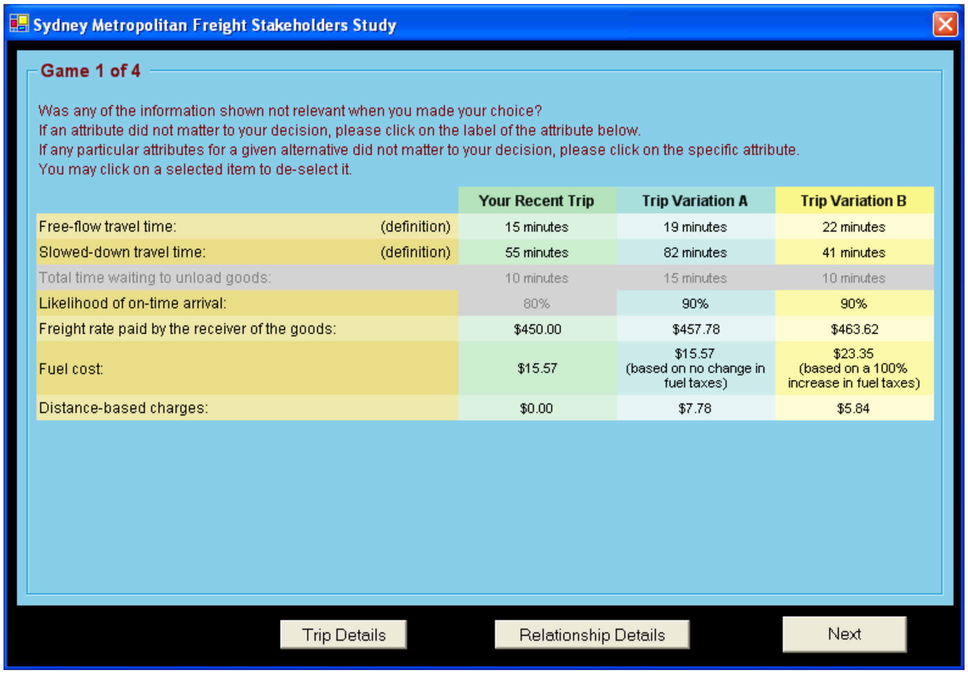

Figure 2: Attribute exclusion screens

Utilising this method, the most complex case of data transformation is the case in which one element of APS choice impacts another. In this empirical exercise, such overlapping is present when a respondent ignores a single element of transit time, allowing the respondent to aggregate a subset of all measures along a common metric. The behavioural meaning of such a nested APS choice is that not all measures along a common metric were important to the process of choosing an alternative, yet, among those that were attended to, each was combined into a composite measure. To represent this particular APS appropriately within the data transformation, each of the steps (1) through (4) above are followed, with the corresponding attribute level for the aggregate specified as the sum of the non-ignored attributes in the aggregate. For example, if a respondent ignored slowed-down time and added the remaining transit time measures together, the data are transformed as follows: (1) a marginal utility of zero is assigned to slowed-down time; (2) marginal utilities of zero are assigned to each transit time measure (making step (1) redundant in this case, but not redundant when other attributes are ignored concurrently); (3) an estimate of the marginal utility of aggregated time is made, with the value (i.e., attribute level) of aggregate time specified as the sum of free-flow time and waiting time; and (4) marginal utility estimates are made for each remaining individual attribute, with a marginal utility of zero assigned to aggregated cost.

A related requirement in models conditioned on APS information centres on the specification of aggregates in the reference alternative. That is, in cases of aggregates that involve attributes that do not appear in the reference alternative, an identification dilemma may arise. The dilemma involves the aggregate of the transporter's costs. The transporter's costs, by definition, are equal to the fuel cost in the reference alternative, due to the absence of distance-based charges in the reference alternative. This does not allow one to simultaneously estimate an individual marginal utility parameter for fuel cost and a marginal utility parameter for aggregated cost within the reference alternative, due to the presence of a singular variance matrix. One potential solution is to specify the cost aggregate only in terms of the stated choice alternatives (i.e., estimate the marginal (dis)utility of aggregated cost only within stated choice alternatives). However, this is not behaviourally meaningful, in that the aggregate measure is intended to reflect the choices of individuals who attend to all costs equally; restricting the marginal utility estimate to choices involving the stated choice alternatives may obscure the resulting inference with respect to this important construct. Indeed, candidate models that included such a specification were poorly behaved, supporting the position that the specification is inappropriate.

Alternatively, one can estimate the marginal utility of the aggregate across all alternatives, whilst estimating the marginal utility of each individual component within the stated choice alternatives. In this exercise, this involves entering the fuel cost as the attribute level for aggregate cost in the reference alternative (i.e., adding the A$0 distance-based charge to the fuel cost to form the effective aggregate). The resulting models discussed in the remainder of this section utilise this method, yielding behaviourally-meaningful results that most closely reflect the attribute aggregation strategies enacted by respondents with respect to transporters' costs. Re-expressing information on APS propensities in terms of alternatives faced, Tables 1 and 2 summarise the degree to which attributes in the model were assigned an adjusted value for marginal utility.[2]

Table 1: Occurrence of attribute exclusion

Attribute	**Number of Times Ignored (Frequency)**
Free-flow time (transporter)	43 (10%)
Slowed-down time (transporter)	53 (12%)
Probability of on-time arrival (transporter)	49 (11%)
Freight rate (transporter)	40 (9%)
Fuel cost (transporter)	29 (7%)
Variable charges (transporter)	25 (6%)

Table 2: Occurrence of attribute aggregation

Attribute	**Number of Times Aggregated (Frequency)**
Time measures (transporter)	292 (68%)
Cost measures (transporter)	326 (75%)

Marginal utilities of zero were assigned in a fairly narrow range across attributes, ranging from a low of six percent for variable charges to a high of 12 percent for slowed-down time. Transporters aggregated all time measures approximately 68 percent of the time, and aggregated all costs approximately 75 percent of the time.

The implications of models without APS information stand to be significantly different to those that incorporate APS information in this application. Whilst this bias may be minor in the case of attributes that are attended to with high frequency, the range of attention paid across attributes is marked in this case. Likewise, the high, but not total, frequency of attribute aggregation could induce biases in models that only allow for a single specification of an attribute that can be added up either on its own or in a composite, but do not allow for both.

[2] We were very careful in establishing a sequence for respondents to work through, and indeed they were able to go back to the previous screen and revise inclusion/exclusion if they did indeed realise subsequently that they meant to include and aggregate specific attributes. Extensive piloting did not raise this concern.

3. Revealing process heterogeneity

Tables 3 and 4 highlight the frequencies with which respondents excluded (i.e., ignored) and aggregated attributes for each alternative in each choice set. The alternatives *RA*, *SC1*, and *SC2* represent the reference alternative, and the first and second stated choice (SC) alternatives, respectively. As above, the choice sets are grouped in the order shown to the respondents:

Table 3: Frequency of attribute exclusion (percent)

Choice Set:	**1**	**2**	**3**	**4**
Free-Flow Time:				
RA	6.5	5.6	7.4	6.5
SC1	6.5	6.5	7.4	8.3
SC2	7.4	6.5	8.3	8.3
Slowed-Down Time:				
RA	8.3	8.3	6.5	7.4
SC1	8.3	8.3	7.4	7.4
SC2	8.3	8.3	7.4	7.4
Likelihood of On-Time Arrival:				
RA	10.2	11.1	12.0	12.0
SC1	10.2	11.1	12.0	12.0
SC2	10.2	11.1	12.0	12.0
Freight Rate:				
RA	10.2	9.3	9.3	8.3
SC1	10.2	9.3	9.3	8.3
SC2	10.2	9.3	9.3	8.3
Fuel Cost:				
RA	5.6	3.7	3.7	4.6
SC1	6.5	3.7	3.7	5.6
SC2	5.6	3.7	5.6	5.6
Variable Charges:				
SC1	3.7	3.7	4.6	3.7
SC2	3.7	3.7	4.6	3.7

The propensity to exclude an attribute varies both across attributes for a given choice set, and across choice sets for a given attribute. Hence, APS heterogeneity plays a strong role, with significant variation across transporters with respect to the manner in which each attribute is processed throughout the choice task. This will be discussed in greater detail below. Table 4 summarises the frequency with which transporters aggregated attributes.

Table 4: Frequency of attribute aggregation (percent)

Choice Set:	**1**	**2**	**3**	**4**
Times:				
RA	65.7	67.6	71.3	72.2
SC1	67.6	69.4	72.2	74.1
SC2	65.7	66.7	73.1	73.1
Costs:				
SC1	76.9	77.8	74.1	76.9
SC2	75.9	77.8	75.0	76.9

The most striking feature of Table 4 is that transporters overwhelmingly chose to aggregate all time and cost measures. A relatively minor proportion chose to treat waiting separate to free-flow and slowed-down time. We observed a range of outcomes with respect to the choice whether to aggregate attributes across choice sets. There was an increase in the propensity to aggregate all time measures as the choice task progresses, offset by a decrease in the propensity to aggregate travel time measures only. Most stable was the choice to aggregate costs, which fluctuated only slightly throughout the choice task for transporters.

3.1 Discussion of processing heterogeneity

There are four primary dimensions along which to examine the degree of variation in each element of APS choice: (1) across attributes; (2) across alternatives within a given choice set; (3) across choice sets for a given alternative; and (4) across classes of decision makers. By analysing the propensity with which attributes were ignored and aggregated along each dimension, one gains a fuller understanding of the choice behaviour and preferences of respondents.

3.1.1 Attribute-specific APS heterogeneity

The degree to which APS choices were enacted for each attribute is the broadest, but most essential element to consider. That is, it is immaterial to consider variations in APS choice across alternatives or choice sets if there is no observed tendency to ignore or aggregate attributes. The stated APS choices of respondents within the sample confirm the importance in capturing information about attribute exclusion and aggregation.

Although some attributes were attended to at high rates by respondents (e.g., cost measures, which were ignored by between three to six percent of transporters, on average), some attributes were ignored by large proportions of the sample. For example, the likelihood of on-time arrival was ignored by approximately 11 percent of transporters.

The results are far more drastic when considering attribute aggregation. Far from being simply an exceptional behaviour that is important to capture, aggregation was prolific within the sample. Transporters aggregated all time measures approximately 70 percent of the time, and aggregated all costs approximately 76 percent of the time.

Hence, the implications of models without APS information stand to be significantly different to those that incorporate APS information. That is, an assumption of full attention to all attributes, each attended to in an equivalent way across decision makers, may induce bias into the estimation process (Puckett and Hensher 2007 confirm this in a mixed logit model). Whilst this bias may be minor in the case of attributes that are attended to with high frequency, the range of attention paid across attributes is marked in this case. Under such circumstances, a passive rationality model in which all attributes are relevant as is may be a dangerous choice. Likewise, the high, but not total, frequency of attribute aggregation could induce biases in models that only allow for a single specification of a malleable attribute (i.e., utility functions that specify an attribute that can be added up either on its own or in a composite, but do not allow for both). The data confirm that APS heterogeneity was present in the sample, supporting the use of modelling structures conditioned on APS information.

3.1.2 Alternative-specific APS heterogeneity

Beginning with cross-alternative variation in APS choice, the data reveal minor variations in the degree to which a particular APS element was enacted. That is, some respondents ignored a particular attribute (or formed a given composite of attributes) in one alternative in a given choice set, while attending to the same attribute (or attending to each element that could form the aggregate) *in another alternative in the same choice set*. As a group, the degree of such variation appears minor, being restricted to a difference of only one or two choices of exclusion per attribute per choice set. When looking at attribute aggregation, the difference becomes more significant, with variations of up to three choices (2.7%).

Viewing the data as a whole could obscure the importance of such variation, however an examination of the frequency with which an individual attribute level was excluded confirms the minor role of alternative specificity in APS choice. The number of times an individual attribute was excluded in the sample ranges from zero (for the likelihood of on-time arrival) to 12 (0.32%). This one empirical study supports the view that there appears no need to specify APS choice at the alternative level.

3.1.3 Choice-set-specific APS heterogeneity

There does appear to be a need to prompt respondents for information regarding their APS choice within each choice set, however. Transporters enacted APSs that were stable across choice sets, with slight changes in the frequency of adopting some APSs. The most significant change in APS for transporters was the decision to aggregate all time measures: in the first choice set faced, transporters aggregated all times between 65.7 percent of the time (for the reference alternative and second SC alternative) and 67.6 percent of the time for the first SC alternative. By the fourth choice set faced, transporters aggregated all times between 72.2 percent of the time (for the reference alternative) and 74.1 percent of the time.

Although the degree to which a given attribute-specific element of respondents' APSs varied across choice sets is not generally severe, the direction of variation is not always toward increased exclusion. That is, respondents appear to have appreciated some attributes more as the choice task progressed, incorporating them into later choices after ignoring them initially. This is counter to the conventional wisdom regarding cognitive burden, which would assume that respondents would either exclude the same attributes across choice sets, or exclude an increasing number of attributes across choice sets as the choice task progresses, due to fatigue effects.

Furthermore, the observation of exclusion propensities increasing across choice tasks for some attributes, whilst decreasing or holding steady for others, supports the notion that respondents took the APS indication task seriously. That is, respondents took the effort to indicate which attributes were ignored, and were willing to indicate when their exclusion strategies changed from choice set to choice set. If such effort were not made, either no fluctuation or non-systematic fluctuation in exclusion propensities may have been observed across choice sets.

Ultimately these results confirm that a choice-set-specific specification of APS prompts is an improvement over one that prompts respondents only upon completing all choice sets. However, the lack of sizeable fluctuations in APS choice across choice sets in either group is insufficient to discredit the practice of specifying APS choice only after the entire choice task has been completed (although there is sufficient gain in choice-set specific variation in model estimation). Whether the information added from a choice-set-specific APS question format becomes more important as either the number of choice sets faced grows, or under different choice settings (e.g., consumer goods purchases, environmental preferences) is an issue for future research.

4. Structural sources of attribute aggregation strategies

In a model of transporters' preferences conditioned on APS information, the choice whether to aggregate times, costs or both leads to highly different behavioural implications. Rather than assume that the sources of aggregation strategies are randomly unobserved, it is important to search for systematic influences on aggregation strategies. This section identifies links between covariates, personal characteristics, relationship characteristics and physical characteristics of the trip, and attribute processing strategies, within an error components logit model framework that accounts for the possibility of differential variance in unobserved sources of influences that are alternative-specific.

The choice set comprises four alternatives: (i) times and costs are aggregated, (ii) times only are aggregated, (iii) costs only are aggregated, and (iv) no attributes are aggregated. The distribution of attribute aggregation strategies across alternatives is shown in Table 5. To assess whether respondents aggregation strategies are influenced by the order with which each choices set is reviewed, we introduce heteroescedastic error components as a way of establishing whether there is any systematic variation in the error components that can be linked to one or more of the four choice sets that each respondent assessed.

Table 5: Frequency of attribute aggregation strategy

Attributes Aggregated	Frequency of Aggregation Strategy
Times and Costs	64.4%
Times Only	5.6%
Costs Only	18.4%
None	11.7%

To set out the error components logit model, we begin with the basic form of the multinomial logit model, with alternative specific constants α_{ji} and attributes x_{ji}, for individuals i = 1,...,N in choice setting t

$$\text{Prob}(y_{it} = j_t) \quad = \frac{\exp\left(\alpha_j + \boldsymbol{\beta}'\mathbf{x}_{jit}\right)}{\sum_{q=1}^{J_i} \exp\left(\alpha_q + \boldsymbol{\beta}'\mathbf{x}_{qit}\right)} \tag{1}$$

where β_k is the population mean for the k^{th} attribute (k=1,…,K) (Train 2003, Hensher *et al.* 2005), and α_j are the choice specific constants. A layer of individual heterogeneity is added to the model in the form of the error components that capture influences that are related to alternatives in contrast to attributes. We do this by constructing a set of independent individual terms, E_{im}, m = 1,...,M ~ N[0,1] that can be added to the utility functions. This device allows us to create what amounts to a random effects model and, in addition, a very general type of nesting of alternatives. Let θ_m be the scale parameter (standard deviation) associated with these effects. Then, each utility function can be constructed as

$$U_{ijt} = \alpha_{ji} + \beta_j' x_{jit} + (\text{any of } \theta_1 E_{i1}, \theta_2 E_{i2}, ..., \theta_M E_{iM}) \tag{2}$$

Consider, for example, a four outcome structure

$$\begin{aligned}
U_{i1t} &= V_{i1t} + \theta_1 E_{i1} + \theta_2 E_{i2} \\
U_{i2t} &= V_{i2t} + \theta_2 E_{i2} \\
U_{i3t} &= V_{i3t} + \theta_1 E_{i1} + \theta_3 E_{i3} \\
U_{i4t} &= V_{i4t} + \theta_4 E_{i4}
\end{aligned}$$

Thus, U_{i4t} has its own uncorrelated effect, but there is a correlation between U_{i1t} and U_{i2t} and between U_{i1t} and U_{i3t}. This example is fully populated, so the covariance matrix is block diagonal with the first three freely correlated. The model might usefully be restricted in a specific application. A convenient way to allow different structures is to introduce the binary variables d_{jm} = 1 if random term E_m appears in utility function j and zero otherwise. Then, the entire model can be specified as equation 3.

$$\text{Prob}(y_{it} = j) = \frac{\exp\left[\alpha_j + \beta' x_{jit} + \Sigma_{m=1}^{M} d_{jm}\theta_m E_{im}\right]}{\sum_{q=1}^{J_i} \exp\left[\alpha_q + \beta' x_{qit} + \Sigma_{m=1}^{M} d_{qm}\theta_m E_{im}\right]}. \tag{3}$$

and the corresponding choice specific dummy variables in x_{jit}. The conditional choice probabilities are defined in (4) (based on Greene and Hensher 2007).

$$\text{Prob}(y_{it} = j|E_i) = \frac{\exp\left[\beta x_{jit} + \Sigma_{m=1}^{M} d_{jm}\theta_m E_{im}\right]}{\sum_{q=1}^{J_i} \exp\left[\beta x_{qit} + \Sigma_{m=1}^{M} d_{qm}\theta_m E_{im}\right]}. \tag{4}$$

The unconditional choice probabilities are formed by integrating the heterogeneity out of the conditional ones. Thus,

$$\text{Prob}(y_{it}=j) = \int_{E_i} \text{Pr}ob(y_{it} = j | E_i) f(E_i) dE_i \tag{5}$$

The integral does not exist in closed form, so we have approximated it with simulation (see Bhat (2003), Revelt and Train (1998), Train (2003) and Brownstone *et al.* (2000) for details). The simulated log likelihood for n individuals and T_i choices by each individual is then

$$\ln L_i^S = \sum_{i=1}^{n} \ln\left\{\frac{1}{R}\sum_{r=1}^{R}\prod_{t=1}^{T_i} \frac{\exp\left[\beta x_{jit} + \Sigma_{m=1}^{M} d_{jm}\theta_m E_{im,r}\right]}{\sum_{q=1}^{J_i} \exp\left[\beta x_{qit} + \Sigma_{m=1}^{M} d_{qm}\theta_m E_{im,r}\right]}\right\}. \tag{6}$$

This model with error components for each alternative is identified. We are estimating the θ parameters as if they were weights on attributes, not scales on disturbances. The parameters are identified in the same way that the βs on the attributes are identified. The parameter on the error component is ($\delta_m\sigma_m$), where σ_m is the standard deviation. Since the scale is not identified, we normalize it to one for estimation purposes, with the understanding that the sign and magnitude of the weights on the error components are carried by θ. The sign of δ_m is also not identified, since the same set of model results will emerge if the sign of every draw on the component were reversed – the estimator of δ would simply change sign with them. Hence we can normalize the sign to plus, and estimate $|\delta_m|$, with the sign and the value of σ_m also normalized for identification purposes.

The final empirical model of choice of attribute aggregation strategy is summarised in Table 6.

Table 6: Error component logit model for four attribute aggregation processing rules parameter (t-statistic)

	Aggregate Times and Costs	**Aggregate Times**	**Aggregate Costs**	**No Aggregation**
Characteristics:				
Years Working in a Similar Role				-0.121 (-2.61)
Years Working in One's Organisation			0.1013 (2.87)	-0.169 (-2.34)
Length of the Business Relationship			-0.124 (-2.66)	
Time to Prepare Truck for Delivery		0.004 (3.93)		
Trip Originated within Urban Area		-1.655 (-3.22)		
Number of Trucks		0.0326 (4.43)		
Number of Delivery Locations	-0.1745 (-2.14)			
Transporter Responsible for Routing			-3.195 (-2.8)	-3.811 (-2.8)
Transporter Responsible for Scheduling	3.3607 (3.06)			
Sender Responsible for Scheduling		-1.637 (-3.40)		
Sender Responsible for routing			-4.733 (-3.5)	
Receiver Responsible for routing			-5.361 (-3.1)	
Time Remaining to Satisfy Delivery	-.0006 (-4.08)		-0.006 (-5.7)	
Kilometres Travelled	0.0060 (2.88)			
Alternative-Specific Constant	-3.413 (-2.44)	-4.448 (-3.67)	1.100 (1.19)	
Standard deviation of latent random effects:	3.052 (2.38)			4.473 (2.94)
Heterogeneity in variance of latent random effects:				
Choice Set #1	0.3639 (1.81)			
Choice Set #2	0.3909 (2.06)			
No. of Observations	1296			
LL(B)	-1038.61			

We will begin the analysis with the most-common attribute aggregation strategy, the choice to aggregate times and costs. With respect to trip characteristics, respondents tend to aggregate times and costs when there were relatively few delivery locations, when the time remaining to satisfy the delivery is relatively short, and when the trip distance is relatively long. Hence, trips involving time pressure and simple routing, especially those containing a large proportion of travel outside of metropolitan areas, are likely to lead to the aggregation of times and costs. An urgent line haul trip of a single commodity to a single destination would be a quintessential example of such a case. Transporters who are responsible for scheduling of trips, are also more likely to aggregate times and costs. That is, transporters who control when trips take place are less sensitive to the mix of free-flow and slowed-down travel time that is experienced during the trip.

Systematic variation was found in the error component for the alternative to aggregate times and costs. Importantly, this variation is a function of how far the choice task has progressed. We find that respondents behave significantly differently within the first two choice sets when choosing whether to aggregate times and costs. Hence, respondents' propensities to enact the most popular, and most extreme, attribute aggregation strategy vary systematically with the amount of information they have processed within the choice task. This is intuitive, in that respondents may learn how best to process the information presented to them in the early stages of a choice task. Importantly we were not able to find any systematic differences in preference heterogeneity for each of the other three APS rules attributable to the sequence of choice sets assessed; which is an encouraging result.

Transporters are more likely to choose to aggregate times, but not costs, when a relatively long time is needed to prepare the truck for delivery, when trips originate from outside of urban areas, when their firms control a large number of trucks, and when the sender does not have responsibility for routing. Hence, transporters find it useful to distinguish between fuel cost and variable road user charges, whilst treating all time measures equally, for complex loads and trips entering metropolitan areas (i.e., areas where road user charges are more likely to be applicable when delivering goods). Likewise, transporters who aggregate time measures are more likely to distinguish between variable charges and fuel cost when they enjoy greater market share and independence with respect to routing of vehicles.

Characteristics of the relationship between transporters and their customers have a strong impact on the decision whether to aggregate cost measures whilst distinguishing between time measures. The longer the two firms have been working together, the less likely transporters are to adopt this attribute aggregation strategy. The more firms with input into the routing of the vehicle, the less likely transporters are to aggregate costs alone. Input from receivers of the goods causes the greatest disutility in aggregating costs alone, followed by input from senders of the goods. The longer the respondent has been working for the transport firm, the more likely the respondent is to aggregate costs alone. Lastly, as with the decision to aggregate both times and costs, transporters are more likely to aggregate costs alone if there is relatively less time available to satisfy the delivery.

Lastly, those who chose not to aggregate any measures tended to have less experience both in their position and with their organisation. Furthermore, those who did not have input into the routing of the vehicle were more likely not to aggregate any measures. Hence, there appears to be a significant link between the propensity to aggregate at least some measures and both experience and influence in the group decision-making process. Those with more experience and influence may be more capable of, or confident in, identifying which of the array of information in a decision-making setting can be grouped together with similar information to make an informed choice more efficiently.

The direct implication of these results is that contextual effects influence preferences, and hence behaviour, through the manner in which the alternatives are evaluated. Whether this is a real-market phenomenon would need to be tested, but the implications are clear within the experimental setting: although contextual effects did not directly explain preference heterogeneity (as reported in Puckett and Hensher 2007), these effects impacted the choice of APS, which is itself a mechanism for explaining preference heterogeneity. Hence, the influence of contextual effects on APS choice equates to an impact of contextual effects on the underlying observed preferences of respondents. That is, contextual effects linked to APS choice are also linked to the behaviour correlated with being within the subset making a particular APS choice.

An important implication of the simultaneous presence of links between APS choice and both contextual effects and choice behaviour is that APS choice may be a more powerful means by which to segment samples than socio-demographic characteristics. That is, whilst the state of practice involves explaining heterogeneity in preferences via contextual effects, this may not be the ideal means by which to explain structural variation in preferences across a sample. Rather, the observed APS choices of respondents may offer a stronger proxy for otherwise unobserved sources of preference heterogeneity. Ultimately, APS heterogeneity represents rich information with respect to the underlying preferences of respondents, which may offer a bridge between the contextual effects correlated with preferences and the decision-making behaviour of respondents that is driven by these preferences.

5. Conclusions

The series of models reveal structural relationships between the attribute processing strategies enacted by respondents and personal, firm-specific, relationship-specific and trip-specific contextual effects governing the choices made. Although this information may be interesting either in itself or as a link to further studies of APS choice, its main role is to identify systematic forces leading to the segmentation of the sample into APS sub-groups, each of which holds a distinct behavioural response to a specific circumstance. That is, in models of independent preferences conditioned on APSs, the inferred behaviour of respondents is distinct across classifications of joint attribute aggregation choice. Hence, it is imperative to identify factors that may lead to divergent behaviour and marginal rates of substitution, especially when these factors may not otherwise reveal a direct link to variations in behaviour.

Not only do these models demonstrate the importance of attribute aggregation in the decision-making process, but the degree of heterogeneity in attribute exclusion strategies across respondents and choice sets also underscore the potential impacts of APS behaviour within stated choice experiments.

Puckett and Hensher (2007) show that process heterogeneity, if behaviourally plausible, does make a difference on very important policy outputs. For example, the impact of the internalisation of APS information into the model on the estimation of value of travel time savings (VTTS) at the broadest level, using the data source herein, is summarised in Table 7. Under the assumption that all information matters fully to all respondents (i.e., passive bounded rationality), all that is needed to form this measure is to take an average of VTTS across types of travel time, weighted by the proportion of the types of travel time. In models conditioned on APS information, an additional step is required, through taking a weighted average across attribute processing strategies. When this calculation is performed, it is revealed that the non-APS-centred model overstates the weighted average VTTS of transporters by a factor of about nine percent, relative to the APS-centred model. That is, when assuming that all attributes are attended to identically across all respondents, alternatives and choice sets, the resulting VTTS is A$50.95 per hour, compared to A$44.36 when conditioning the model on the available APS information.

Table 7: VTTS Measures (Puckett and Hensher 2007)

Weighted average across free-flow & slowed-down time & aggregation strategies at the mean, A$ per hr

	Weighted Average VTTS
Without APS Heterogeneity	$50.95
With APS Heterogeneity	$44.36

This over-estimation when ignoring process heterogeneity has been found to occur in other studies where the process rules are informed only at the completion of all choice scenarios (Hensher *et al.* 2006, Hensher and Rose 2005, Hensher 2008). There is a clear trend emerging on upward biased mean estimate of marginal willingness to pay when ignoring process heterogeneity.

6. References

Bonini, N., Tentori, K. and Rumiati, R. (2004) Contingent application of the cancellation editing operation: the role of semantic relatedness between risky outcomes, *Journal of Behavioral Decision Making*, 17, 139-152.

Hensher, D.A. (2006) How do respondents handle stated choice experiments? - attribute processing strategies under varying Information load, *Journal of Applied Econometrics*, 21, 861-78.

Hensher, D.A. (2008) Joint estimation of process and outcome in choice experiments and implications for willingness to pay, *Journal of Transport Economics and Policy*.

Hensher, D.A. and Rose, J. (2005) Respondent Behavior in Discrete Choice Modeling with a Focus on the Valuation of Travel Time Savings, *Journal of Transportation and Statistics,* 8 (2),17-30.

Hensher, D.A., Rose, J. and Greene, W. H. (2005) *Applied Choice Analysis: A Primer,* Cambridge University Press, Cambridge.

Hensher, D.A., Rose, J. and Bertoia, T. (2006) The implications on willingness to pay of a stochastic treatment of attribute processing in stated choice studies, *Transportation Research E*, 43, 73-89.

Hensher, D.A., Puckett, S. and Rose, J. (2007) Agency decision making in freight distribution chains: revealing a parsimonious empirical strategy from alternative behavioural structures, *Transportation Research B*. 41 (9), 924-949.

Hess, S., Rose, J. M. and Hensher, D.A. (2008) Asymmetrical preference formation in willingness to pay estimates in discrete choice models (presented only at ETC2006 Strasbourg and TRB 2007), *Transportation Research* E

Kahneman, D. and Tversky, A. (1979) Intuitive prediction: biases and corrective procedures. In S. Makridakis and S. C. Wheelwright, Eds., *Studies in the Management Sciences: Forecasting*, 12 (Amsterdam: North Holland).

Puckett, S.M. and Hensher, D.A. (2007 in press) The role of attribute processing strategies in estimating the preferences of road freight stakeholders under variable road user charges, *Transportation Research E*.

Rabin, M. (1998) Psychology and economics, *Journal of Economic Literature*, XXXV1, March, 11046.

Revelt, D. and K. Train (1998) Mixed Logit with repeated choices: households' choices of appliance efficiency level, *Review of Economics and Statistics*, 80, 1-11.

Rose, J. and Bliemer, M. (2007) Stated preference experimental design stratgegies in Hensher, D.A, and Button, K.J. (Eds.), *Handbook of Transport Modelling*, second edition, Elsevier, Oxford,151-180.

Train, K.E. (2003) *Discrete Choice Methods with Simulation*, Cambridge University Press, Cambridge.