Technical Appendix – Description of Welfare Measurement Algorithm and Additional Results

(not intended for publication – available upon request)

1. Welfare Measurement

1a. Generic Issues

The approach used for constructing Hicksian welfare measures builds on the strategy developed by von Haefen, Phaneuf and Parsons (2004) (vHPP hereafter). They develop a multi-stage, Monte Carlo, Markov Chain simulation algorithm that solves for the expected compensating surplus arising from changes in price and/or quality. Since their algorithm applies to situations where the consideration set is known and fixed, it must be extended to the present setting where the consideration set is latent and probabilistic.

Using standard notation, the Hicksian consumer surplus, CS^H, associated with a price and quality change from ($\boldsymbol{p}^0, \boldsymbol{Q}^0$) to ($\boldsymbol{p}^1, \boldsymbol{Q}^1$) is implicitly defined as:

(A1) $$V(\boldsymbol{p}^0, \boldsymbol{Q}^0, y, \boldsymbol{\beta}, \boldsymbol{\varepsilon}) = V(\boldsymbol{p}^1, \boldsymbol{Q}^1, y - CS^H, \boldsymbol{\beta}, \boldsymbol{\varepsilon}).$$

vHPP identify three related issues that complicate recovering CS^H from traditional continuous demand system models. Unless preferences are homothetic or quasilinear in income, no closed form solution for the income reduction that equates utility across the two states exist, and thus iterative search procedures such as numerical bisection must be employed to solve for CS^H. Moreover, at each iteration of the search procedure, the analyst must solve for the consumer's utility conditional on a $(\boldsymbol{p}, \boldsymbol{Q}, y, \boldsymbol{\varepsilon})$ quadruplet. As vHPP argue, solving the consumer's problem analytically as in Phaneuf, Kling and Herriges (2000) quickly becomes computationally intractable as the dimension of the consideration set grows large, and thus numerical approaches must be used for this task. Finally, because $\boldsymbol{\varepsilon}$ is a random variable from the analyst's perspective, she cannot ascertain CS^H precisely. At best the analyst can recover an estimate of the central tendency of CS^H such as its expectation, $E(CS^H)$, which often requires simulation to construct.

All of these difficulties are present with models that incorporate latent consideration sets, but the difficulties associated with solving the consumer's problem conditional on a $(\boldsymbol{p}, \boldsymbol{Q}, y, \boldsymbol{\varepsilon})$ quadruplet are more complex because there are 3^M regimes that the individual might choose compared to 2^M for traditional models. As a result, the analyst must first determine which goods the individual considers and then conditionally the combination of interior and corner solutions from her consideration set that represents her optimal consumption bundle. Because the structure of the consideration process implies that the individual's consideration set is a function of prices, quality, and income, the analyst must resolve for the individual's consideration set whenever these values change. This fact suggests that the analyst must resolve for the individual's consideration set at every iteration of the search routine that solves for CS^H. Thus, the computationally challenge of welfare measurement arising from models with latent consideration sets is more formidable relative to traditional demand system models.

Similar to vHPP, the simulation-based approach to constructing Hicksian consumer surplus measures can be decomposed into three stages. Although the third stage is trivial (averaging the simulated welfare estimates), the first two are not and therefore discussed in detail in the next two sections.

1b. First Stage – Simulating the Unobserved Heterogeneity Entering Consumer Preferences

Following von Haefen (2003), this paper employs an approach to constructing welfare measures sets that incorporates the implications of an individual's observed choice. In other words, $\boldsymbol{\varepsilon}$ is simulated such that the model perfectly predicts observed behavior at baseline conditions, and the structure of substitution implied by the model is used to predict how people respond to price, quality, and income changes. In the context of traditional demand system models, vHPP find that fewer simulations are necessary to achieve an arbitrary level of precision in the welfare estimates relative to more traditional welfare construction approaches that do not incorporate observed choice. These computational advantages carry over to the present context where consideration sets are latent and probabilistic.

Implementation of this approach requires that the analyst simulate from the joint distribution of the unobserved heterogeneity conditional on the individual's observed choice, i.e., $f(\boldsymbol{\varepsilon} \mid \boldsymbol{x}^*)$. The following decomposition suggests a convenient strategy for accomplishing this task:

(A2) $$f(\boldsymbol{\varepsilon} \mid \boldsymbol{x}^*) = f_1(\varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda \mid \boldsymbol{x}^*)\, f_2(\bar{\boldsymbol{\varepsilon}}' \mid \varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda, \boldsymbol{x}^*)\, f_3(\bar{\boldsymbol{\varepsilon}} \mid \bar{\boldsymbol{\varepsilon}}', \varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda, \boldsymbol{x}^*)\,.$$

Equation (A2) states that the joint distribution of the unobserved heterogeneity conditional on $\boldsymbol{x}^*$ can be decomposed into the marginal distribution of the random coefficients multiplied by the distribution of $\bar{\boldsymbol{\varepsilon}}'$ conditional on the random coefficients and $\boldsymbol{x}^*$ and the distribution of $\bar{\boldsymbol{\varepsilon}}$ conditional on $\bar{\boldsymbol{\varepsilon}}'$, the random coefficients, and $\boldsymbol{x}^*$. This structure suggests a sequential simulation procedure where the analyst first simulates the random coefficients from $f_1(\varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda \mid \boldsymbol{x}^*)$ and then conditionally simulates $\bar{\boldsymbol{\varepsilon}}'$ from $f_2(\bar{\boldsymbol{\varepsilon}}' \mid \varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda, \boldsymbol{x}^*)$, and finally $\bar{\boldsymbol{\varepsilon}}$ from $f_3(\bar{\boldsymbol{\varepsilon}} \mid \bar{\boldsymbol{\varepsilon}}', \varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda, \boldsymbol{x}^*)$.

As discussed in von Haefen and vHPP, simulating from $f_1(\varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda \mid \boldsymbol{x}^*)$ can be accomplished by exploiting a Metropolis-Hastings simulation algorithm (Chib and Greenberg, 1995). Since the procedures involved in implementing the Metropolis-Hastings algorithm are fairly standard and outlined by vHPP in the context of traditional Kuhn-Tucker demand system models, the details of the algorithm tailored to the current application are not reported here. It should be noted, however, that because the Metropolis-Hastings algorithm, like all Monte Carlo, Markov Chain simulators, generates simulations from the target distribution only after a sufficiently long burn-in period, the analyst should discard the first *T* simulations. Moreover, to reduce the serial correlation in the Markov Chain of random parameters, the analyst should only use every *j*th simulation ($j > 1$) after the first *T* in subsequent steps.

The second step, however, involves complexities not fully addressed elsewhere. The approach for simulating $\bar{\boldsymbol{\varepsilon}}'$ pursued in this paper involves a sequential procedure that can be interpreted as an extension to an approach suggested originally by von Haefen and Phaneuf (2003) for determining whether a non-consumer of a set of quality differentiated goods would consider consuming any of the goods under any circumstance. A feature of the approach used here is that it does not simulate $\bar{\boldsymbol{\varepsilon}}'$ directly but instead simulates the individual's consideration set. Given the structure of the consideration process, it should be clear that simulating $\bar{\boldsymbol{\varepsilon}}'$ and simulating the individual's consideration set are observationally equivalent conditional on prices, quality, and income.

The initial step in this process is to simulate the individual's consideration set under baseline conditions when the analyst observes a choice. Because commodities must be considered to be consumed, the analyst knows with certainty that all goods with interior solutions (i.e., strictly positive demands) are considered by the individual at baseline conditions. For goods with corner solutions, the analyst cannot identify precisely whether they are considered. However, the structure of the econometric model can inform the analyst about the likelihood of consideration conditional on observing a corner solution. In particular, the probability of an individual considering x_j^* when it is not consumed conditional on $(\varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda)$ is:

$$\text{(A3)} \qquad \Pr(x_j^* \text{ considered}|\varepsilon_\delta,\varepsilon_\kappa,\varepsilon_\lambda) = \frac{\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)l_j(\varepsilon_\delta \mid x_j^*=0)}{(1-\pi_j(\varepsilon_\kappa,\varepsilon_\lambda))+\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)l_j(\varepsilon_\delta \mid x_j^*=0)},$$

where $\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)$ and $l_j(\varepsilon_\delta \mid x_j^*=0)$ are evaluated at baseline prices, quality, and income. Simulating whether x_j is considered by the individual can be accomplished by comparing a uniform random draw, *U*, to (A3), (i.e., x_j^* is considered if $U < \Pr(x_j^* \text{ considered}|\varepsilon_\delta,\varepsilon_\kappa,\varepsilon_\lambda))$.

Once the analyst has simulated the individual's consideration set at baseline conditions, she can then determine how the choice set changes in response to price, quality, and income changes. To see how this is accomplished, it is important to recognize that from a statistical perspective, the structure of the consideration set process implies that x_j is considered by the individual if a uniform random draw, say η_j, is less than the threshold $\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)$. Thus, the analyst knows that if x_j is considered at baseline conditions, $\eta_j < \pi_j(\varepsilon_\kappa,\varepsilon_\lambda)$ when $\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)$ is evaluated at $(\boldsymbol{p}^0,\boldsymbol{Q}^0,y)$. By simulating $\boldsymbol{\eta} = [\eta_1,...,\eta_M]^{\mathrm{T}}$ such that it is consistent with the simulated consideration set at baseline conditions, the analyst can ascertain how the choice set changes for any $(\boldsymbol{p},\boldsymbol{Q},y)$ triplet by comparing the simulated $\boldsymbol{\eta}$ vector with the corresponding $\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)$ probabilities evaluated at any $(\boldsymbol{p},\boldsymbol{Q},y)$.

Operationally, the approach for ascertaining how the simulated choice set changes involves the following steps. Conditional on a simulated consideration set at baseline conditions, the analyst simulates the separate elements of $\boldsymbol{\eta}$ according to the following rule:

$$\text{(A4)} \qquad \eta_j = \begin{cases} \pi_j(\varepsilon_\kappa,\varepsilon_\lambda)U & \text{if } x_j \text{ is considered} \\ 1-\pi_j(\varepsilon_\kappa,\varepsilon_\lambda)U & \text{if } x_j \text{ is not considered} \end{cases},$$

where *U* again is a uniform random draw and $\pi_j(\varepsilon_\kappa, \varepsilon_\lambda)$ is evaluated at $(\boldsymbol{p}^0, \boldsymbol{Q}^0, y)$. Given the simulated $\boldsymbol{\eta}$, the analyst need only compare each η_j to its corresponding $\pi_j(\varepsilon_\kappa, \varepsilon_\lambda)$ evaluated at any $(\boldsymbol{p}, \boldsymbol{Q}, y)$ triplet to determine how the individual's choice set has changed.

The third and final step for simulating the unobserved heterogeneity involves simulating $\bar{\boldsymbol{\varepsilon}}$ from $f_3(\bar{\boldsymbol{\varepsilon}} \mid \bar{\boldsymbol{\varepsilon}}', \varepsilon_\delta, \varepsilon_\kappa, \varepsilon_\lambda, \mathbf{x}^*)$. Assuming good *j* is consumed in strictly positive quantities, the structure of the behavioral model implies that $\varepsilon_j = g_j(\varepsilon_\delta)$ where $g_j(\varepsilon_\delta)$ is defined in equation (8). When x_j is not consumed but considered by the individual, then ε_j must be strictly less than $g_j(\varepsilon_\delta)$. Given the i.i.d. Type I Extreme Value distributional assumption, ε_j can be simulated via

(A5) $$\varepsilon_j = -\ln(-\ln(\exp(-\exp(-g_j(\varepsilon_\delta)/\mu))U))\mu,$$

where μ is the scale parameter and *U* is another random draw from the uniform distribution. If x_j is not consumed or considered, the analyst infers nothing about the values that ε_j can take. In this case, ε_j is simulated from the Type I Extreme Value distribution with full support:

(A6) $$\varepsilon_j = -\ln(-\ln U)\mu.$$

1c. Second Stage – Solving for the Simulated Hicksian Consumer Surplus

The approach for solving for the Hicksian consumer surplus associated with a price and quality change used in this paper employs the same multi-layered numerical search routine employed by vHPP. At the top layer, a numerical bisection routine is used to solve for the income reduction that equates utility before and after the price and quality change. This task is complicated by the fact that the individual's consideration set can change when prices, quality, and income change. These changes in an individual's consideration set generate discontinuities in utility, and thus an income reduction that exactly equates utility before and after a price and quality change might not exist.

Figure A1 suggests the nature of the problem as well as a plausible resolution. Before the quality change, the individual achieves a utility level equal to $V(\boldsymbol{p}^0, \boldsymbol{Q}^0, y, \boldsymbol{\beta}, \boldsymbol{\varepsilon})$. Holding income constant, assume the individual is made better off with the price and quality change to ($\boldsymbol{p}^1, \boldsymbol{Q}^1$), and so an income reduction is necessary to equate utility before and after the price and quality change. Initially as income is taken away from the individual, her utility continuously declines. However, once the individual's adjusted income equals $\bar{y}$, any further reduction generates a discontinuous and potentially large drop in utility. This discontinuity arises because $\bar{y}$ represents a threshold at which the individual's consideration set contracts in a way that a good that is consumed in positive quantities at income levels at and above $\bar{y}$ is not considered and consumed at income levels below it. Figure A1 is drawn such that this discontinuity falls over a utility range that encompasses the baseline utility level, $V(\boldsymbol{p}^0, \boldsymbol{Q}^0, y, \boldsymbol{\beta}, \boldsymbol{\varepsilon})$. Thus, there is no income compensation that equates utility before and after the price and quality change in this example. However, Figure 1 suggests that if one defines CS^H as:

(A7)
$$\begin{array}{l} V(\boldsymbol{p}^0, \boldsymbol{Q}^0, y, \boldsymbol{\beta}, \boldsymbol{\varepsilon}) \le V(\boldsymbol{p}^1, \boldsymbol{Q}^1, y - CS^H + \Delta, \boldsymbol{\beta}, \boldsymbol{\varepsilon}) \\ V(\boldsymbol{p}^0, \boldsymbol{Q}^0, y, \boldsymbol{\beta}, \boldsymbol{\varepsilon}) \ge V(\boldsymbol{p}^1, \boldsymbol{Q}^1, y - CS^H - \Delta, \boldsymbol{\beta}, \boldsymbol{\varepsilon}) \end{array},$$

for any $\Delta > 0$, a unique welfare measure exists. This definition of the Hicksian consumer surplus is used in this paper.

To recover welfare estimates based on (A7), a second numerical bisection routine that solves for the individual's utility conditional on a simulated consideration set and $(\boldsymbol{p}, \boldsymbol{Q}, y, \boldsymbol{\varepsilon})$ quadruplet is necessary. The approach pursued here borrows from vHPP's numerical bisection procedure designed for this same task with the only difference being that the consideration set defines the relevant goods in the current case whereas all goods are relevant in vHPP's setup. As described in vHPP, the routine exploits the fact that when preferences are additively separable, solving for the optimal consumption bundle reduces to a one-dimensional search for the optimal value for *z*, the essential Hicksian composite commodity. The steps of the algorithm are as follows:

1) At iteration *i*, set $z_a^i = (z_l^{i-1} + z_u^{i-1})/2$. To initialize the algorithm, set $z_l^0 = 0$ and $z_u^0 = y$.
2) Conditional on z_a^i, solve for $\boldsymbol{x}^i$ using
$$u_j(x_j) \leq u_z(z) p_j, \forall j \in C$$
$$x_j \geq 0, \forall j \in C$$
$$x_j(u_j(x_j) - u_z(z) p_j) = 0, \forall j \in C$$
where $u_j(x_j)$ and $u_z(z)$ are the additively separable components of the indirect utility function that depend only on x_j and *z*, respectively.
3) Use the budget constraint (i.e., $z = y - \sum_{j \in C} p_j x_j$) and $\boldsymbol{x}^i$ to construct $\tilde{z}^i$.
4) If $\tilde{z}^i > z_a^i$, set $z_l^i = z_a^i$ & $z_u^i = z_u^{i-1}$. Otherwise, set $z_l^i = z_l^{i-1}$ & $z_u^i = z_a^i$.
5) Iterate until $abs(z_l^i - z_u^i) \leq c$ where *c* is arbitrarily small.

Using the same arguments employed by vHPP, it is straightforward to show that the curvature properties of the direct utility function imply that the algorithm will solve for the consumer's optimal consumption bundle to any arbitrarily defined level of precision. Once these values are recovered, the analyst can solve for the individual's utility conditional on a $(\boldsymbol{p}, \boldsymbol{Q}, y, \boldsymbol{\varepsilon})$ quadruplet by plugging them into (6).

1d. Empirical Implementation

Although the algorithm described above has more layers and details than the algorithm developed by vHPP, it remains relatively straightforward to code within a matrix programming language. Experience with the algorithm in an applied setting suggests, however, that significantly more simulations are necessary to achieve an arbitrary level of precision in the welfare point estimates relative to traditional demand system models with fixed and known consideration sets. Nevertheless, constructing consistent point and standard errors estimates for welfare measures from Kuhn-Tucker demand system models for models with latent and potentially large consideration sets remains well within the realm of current computational feasibility.

2. Additional Results

Table A1 reports the correlation parameter estimates from the correlated coefficient models with and without consideration spikes. A number of these estimates are statistically significant at the 95 percent level. Table A2 reports results from consistent Akaike information criteria Schwartz (1978), likelihood ratio tests (Self and Liang, 1987; Andrews, 2001), and Vuong non-nested hypothesis tests (Vuong, 1989) that are used to compare the different models. In general, these results suggest that: 1) the latent consideration set models fit the data better than the traditional models; and 2) the evidence is somewhat mixed with respect to whether the fixed or random parameter models fit the data better.

3. References

Andrews DW. 2001. Testing when a parameter is on the boundary of a maintained hypothesis. *Econometrica* **69**: 683-734.

Chib S, Greenberg E. 1995. Understanding the Metropolis-Hastings Algorithm. *American Statistician* **49**: 327-335.

Phaneuf D, Kling C, Herriges J. 2000. Estimation and Welfare Calculations in a Generalized Corner Solution Model with an Application to Recreation Demand. *Review of Economics & Statistics* **82**: 83-92.

Schwarz G. 1978. Estimating the dimension of a model. *Annals of Statistics* **6**: 461-464.

Self S, Liang K. 1987. Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under non-standard conditions, *Journal of the American Statistical Association*. 82(398): 605-610.

von Haefen R. 2003. Incorporating Observed Choice into the Construction of Welfare Measures from Random Utility Models. *Journal of Environmental Economics & Management* **45**: 145-165.

von Haefen R, Phaneuf D. 2003. Estimating Preferences for Outdoor Recreation: A Comparison of Continuous and Count Data Approaches. *Journal of Environmental Economics & Management* 45: 612-630.

von Haefen R, Phaneuf D, Parsons G. 2004. Estimation and Welfare Analysis with Large Demand Systems. *Journal of Business & Economic Statistics* **22**(2): 194-205.

Figure A1
Welfare Measurement

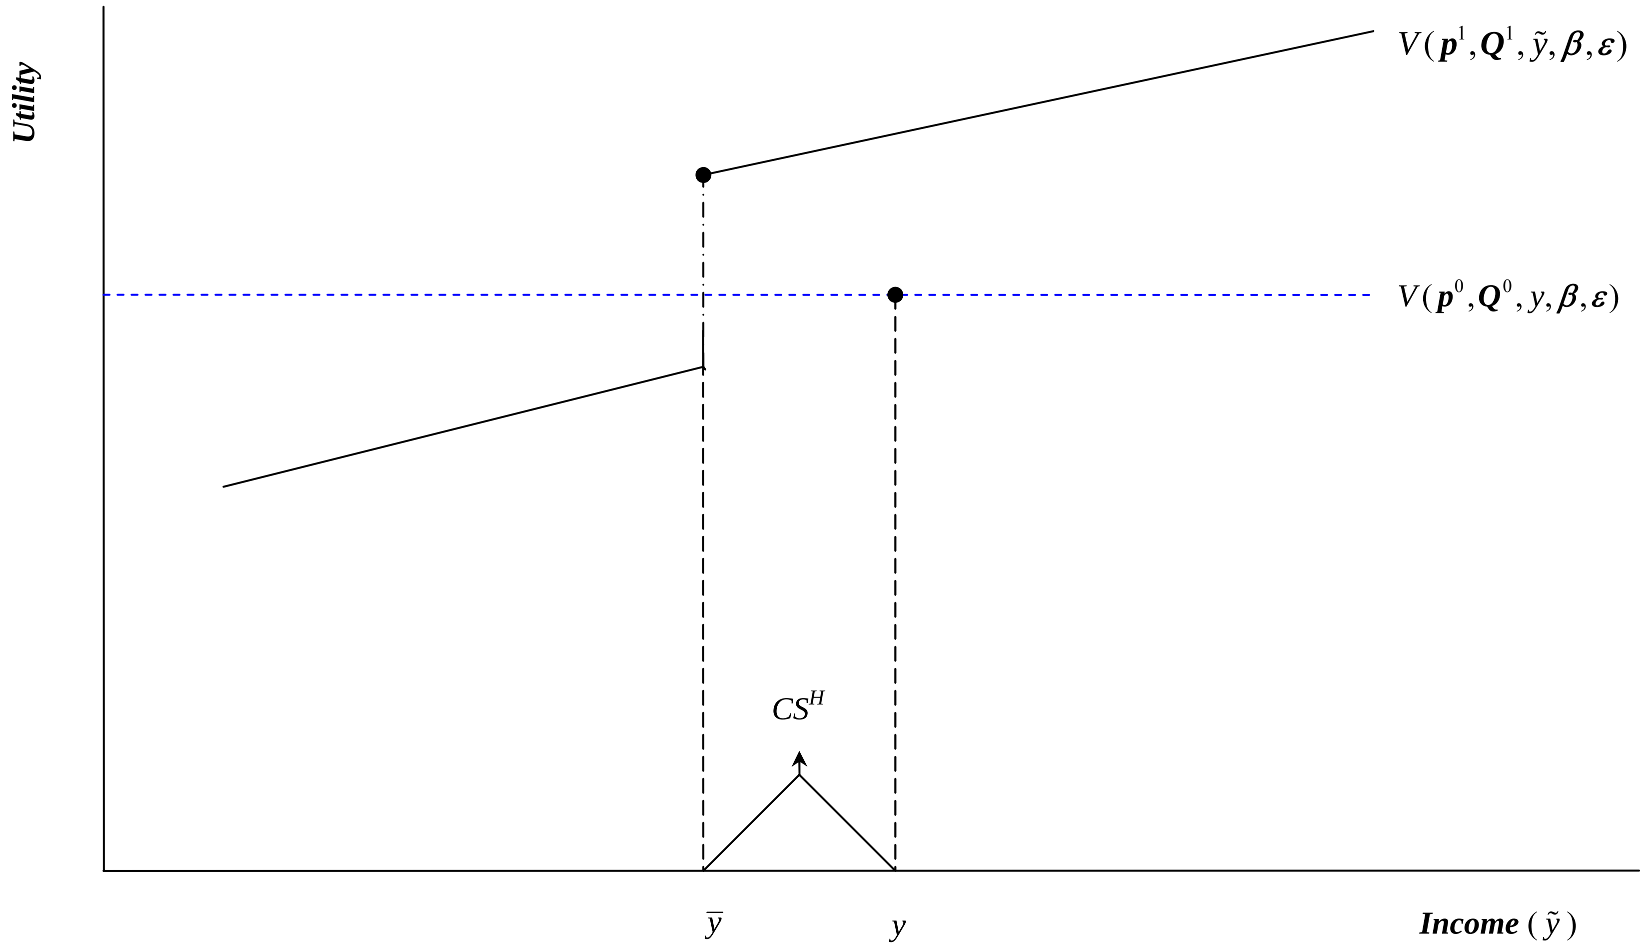

Table A1
Additional Parameter Estimates for Correlated Traditional & Latent Consideration Set Models[a,b]

	Traditional Model Correlated Random Parameters	*Latent Consideration Set Models Correlated Random Parameters*
Parameter correlation between		
$\delta_{WaterRec}$ & δ_{Female}	-0.938	-0.378
	(1.74)	(0.89)
$\delta_{WaterRec}$ & $\zeta_{Susquehanna}$	**-0.750**	**-0.530**
	(4.54)	(3.39)
$\delta_{WaterRec}$ & ζ_{Park}	-0.053	-0.307
	(0.16)	(0.86)
$\delta_{WaterRec}$ & κ	-	-0.276
		(0.90)
$\delta_{WaterRec}$ & $\lambda/10,000$	-	**0.276**
		(2.28)
δ_{Female} & $\zeta_{Susquehanna}$	**0.913**	**0.981**
	(1.96)	(4.38)
δ_{Female} & ζ_{Park}	-0.270	-0.614
	(0.49)	(1.57)
δ_{Female} & κ	-	-0.352
		(1.10)
δ_{Female} & $\lambda/10,000$	-	**0.758**
		(2.55)
$\zeta_{Susquehanna}$ & ζ_{Park}	**-0.620**	-0.541
	(2.69)	(1.52)
$\zeta_{Susquehanna}$ & κ	-	-0.300
		(1.08)
$\zeta_{Susquehanna}$ & $\lambda/10,000$	-	**0.638**
		(5.03)
ζ_{Park} & κ	-	**0.737**
		(3.39)
ζ_{Park} & $\lambda/10,000$	-	**-0.801**
		(2.89)
κ & $\lambda/10,000$	-	**-0.662**
		(2.81)

[a] Random parameter models estimated with 500 Halton draws.
[b] Absolute values of t-statistics based on robust standard errors in parentheses unless otherwise noted. All boldface parameters are significant at the 5% level.

Table A2
Statistical Comparisons of Alternative Models

	Traditional Model - Fixed	***Traditional Model - Uncorrelated***	***Traditional Model - Correlated***	***Latent Consid. Sets Model - Fixed***	***Latent Consid. Sets Model – Uncorrelated***	***Latent Consid. Sets Model - Correlated***
Traditional Model - Fixed	LL(-2,005.84) CAIC(4,072.26)					
Traditional Model – Uncorrelated	LR(←,0.0000)	LL(-1,995.85) CAIC(4,076.48)				
Traditional Model – Correlated	LR(←,0.0002)	LR(←,0.0751)	LL(-1,990.12) CAIC(4,101.36)			
Latent Consid. Sets Model – Fixed	V1(D,0.0000) V2(←,0.0000)	V1(D,0.0000) V2(←,0.0007)	V1(D,0.0000) V2(←,0.0085)	LL(-1,960.22) CAIC(3,999.18)		
Latent Consid. Sets Model – Uncorrelated	V1(D,0.0000) V2(←,0.0000)	V1(D,0.0000) V2(←,0.0000)	V1(D,0.0000) V2(←,0.0001)	LR(←,0.0005)	LL(-1,950.58) CAIC(4,016.24)	
Latent Consid. Sets Model – Correlated	V1(D,0.0000) V2(←,0.0000)	V1(D,0.0000) V2(←,0.0000)	V1(D,0.0000) V2(←,0.0000)	LR(←,0.0016)	LR(↑,0.1067)	LL(-1,939.56) CAIC(4,085.04)

Key

LL(·) & CAIC(·)= Log-Likelihood and Consistent Akaike Information Criteria, respectively.

V1(·,·) = Not statistically distinguishable (ND) or statistically distinguishable (D) in 1st stage of Vuong non-nested hypothesis test (α = .1). P-value also reported.

V2(·,·) = If distinguishable in 1st stage, 2nd stage Vuong test results reported with arrow pointing to statistically prefered model (α = .1). P-value also reported.

LR(·,·) = Likelihood ratio test with arrow pointing to restricted model if it could not be rejected, unrestricted model otherwise (α = .1). P-value also reported.

All likelihood ratio tests employed the standard critical values from the chi-squared distribution with the appropriate degrees of freedom. For the likelihood ratio tests comparing fixed and random coefficient specifications, p-values were constructed using the approach developed by Self and Liang (1987) and Andrews (2001).