WORK IN PROGRESS: DO NOT QUOTE
PAPER PRESENTED AT HESG CONFERENCE, BRUNEL, SEPT 2007

A Model of Adaptive Global Utility Maximization

W. DAVID BRADFORD
Center for Health Economic and Policy Studies, Medical University of South Carolina
Department of Health Administration and Policy, Medical University of South Carolina

PAUL DOLAN
Tanaka Business School, Imperial College London

Abstract

This paper expands the standard model of utility maximization in order to endogenize the ubiquitous phenomenon of adaptation. We consider how the standard utility maximizing model could be expected to adjust to health-related shocks. We assume utility is an aggregate function of the satisfaction associated with domains of life, with weights which are optimized according to the effort that is expended by the individual (using an analog to the household production function). Comparative statics from the general maximization problem demonstrate that the traditional Slutsky equation should incorporate an additional response term. One of the examples we explore is obesity. Our model suggests that there are at least three classes of barriers to controlling obesity: individuals may find it difficult to control their BMI because the opportunity costs of doing so (in direct medical costs or foregone pleasure) are too high; they may have difficulty in translating changes in BMI into changes in actual utility because of the limits in their individual production technology; or they may ultimately fail to reduce BMI because they find it less costly to adapt to lower values of some utility domains. These findings have important implications for service provision.

Corresponding Author:
W. David Bradford, Ph.D., Department of Health Administration, Medical University of South Carolina, 151 Rutledge Avenue - Building B, P.O. Box 250961, Charleston, SC 29425; Phone: (843) 792-2117; Fax: (843) 792-3327; E-mail: bradfowd@musc.edu

The spelling in this paper reflects that used in the US, and we (especially Paul, of course) apologize for any upset this may cause HESG delegates.

1. INTRODUCTION

The canonical model of economic decision making since the time of Jevons, Menger and Walras has been based on the assumption that individuals choose different amounts of goods in order to achieve the maximum possible value of a utility function, given the constraints imposed by personal income and exogenously set prices. This simple model has proved enormously powerful and widely applicable. However, the fundamental simplicity necessarily limits its applicability somewhat. Consequently, through the decades economists have proposed numerous modifications to this basic theory to accommodate decision making under uncertainty (von Neumann and Morgenstern, 1944), the provision of non-market activities (Becker, 1965; Lancaster, 1966), and reference-point behaviors (Tversky and Kahneman, 1991), to name but three of the most prominent. In each case, the models expand upon the underlying assumption that agents' short-term decisions are made in a manner consistent with maximizing a function that adheres to five fundamental axioms of preference (completeness, reflexivity, transitivity, continuity, and non-satiation).

In this paper, we continue this tradition by asking whether the standard model of utility maximization can be redefined in a way that brings a class of the ubiquitous phenomenon of adaptation "into the fold". To appreciate the difficulty with which adaptation can be explained within the standard model, consider an agent who is maximizing utility and receives an exogenous, permanent, increase in real income. The standard model clearly predicts that utility will rise permanently, as all of the increase in income will be allocated completely across the bundle of goods being consumed. Since, by definition, consumption has risen, utility must rise by the full amount of the change in each commodity consumed multiplied by the appropriate marginal utilities. However, one of the more persistent empirical findings on the relationship between income and measures of happiness or well-being is that well-being often does not rise by anything like the full amount expected (Easterlin, 1974; Easterlin, 1995). This dampened response of well-being to changes in income is difficult to explain using a standard neoclassical utility maximization model.

As a second example, consider an agent who is living in a state of constrained utility maximization. She has arranged her budget such that the goods purchased achieve the highest possible level of utility, given her resources, market prices, household production technology, and the like. If this agent is in an automobile accident, and loses the use of her legs, we would expect utility to fall as some of the goods in her current bundle are no longer effective at generating happiness and as her household production technology would no longer be optimized to her new constraints. Whilst this is consistent with the standard model, what follows is not predicted: in most studies of well being, we find that the value of the maximand will rise again to approach (and sometimes equal) its original value, even though the newly restricted resource constraints have not been relaxed (Dolan and Kahneman, 2007).

The problem is not constrained to the relatively manageable issue of how individuals may respond to extreme changes in their income or physical health. A sedentary person may realize that they are compromising their long-term welfare by failing to exercise. One response to this realization may be to attempt to lose weight but the person may find that the opportunity cost per unit of utility improvement is lower from simply migrating toward a peer group that is similarly overweight – thus shifting their revealed preferences away from health and more toward some substitute domain of utility. In other words, it may be easier for the agent to adapt to being overweight than it is to actually reduce their weight.

As a final example, consider how a person who is dissatisfied with their job and who faces a high opportunity cost of switching. He may choose to react to the consequent job-related stress by seeking counseling, which can be viewed as an attempt to develop more realistic career goals thereby changing his revealed tastes with respect to his job. Alternatively, the individual may invest in leisure pursuits that make the utility from work less important – as many people who 'live for the weekend' would do. (Note that this example is very much in the spirit of Becker and Mulligan, who model a person's decision to invest effort toward changing their rate of time preference when she realizes that she is too impatient.)

In this paper, we posit a utility maximizing framework that explicitly incorporates these kinds of adaptation processes. Adaptation is a widely studied phenomenon in psychology, though much less well studied in economics (the word does appear in the index of any core economics textbook that we are aware of). By adaptation, we mean a mechanism that causes the level of utility to change even in the face of constant resource constraints, prices and income. One possible response to losing the use of one's legs may be to invest in prosthetics, modifications to automobiles and the home, and other technological compensations. As a result of these adjustments the agent may be able to increase utility towards the pre-injury level. This is not what we mean by adaptation, which only occurs when utility adjusts independent of any changes to the commodity bundle (this is comparable to the definition found in Menzel et al, 2002).

The remainder of this paper will present a brief discussion of past extensions of utility maximization as well as a more detailed exploration of the literature on adaptation that is relevant to our framework. Section 3 follows with a formal development of our adaptive utility model, including demonstrating that traditional Slutsky-type equations follow from the comparative statics. Generally, preferences are taken as fundamental – the fixed points around which our existence theorems turn – and so solutions that posit changes in preferences are unattractive. Thus, we need a model that maintains the constancy of preferences, and yet allows adjustments to the nature of the utility function such that it can respond in an adaptive manner to external shocks. Finally, Section 4 concludes with a discussion of the empirical implications of the model and of several potential applications.

2. BACKGROUND

2.1. Utility Maximization as Household Production:

The main modification to the standard model is that individuals maximize utility across life domains. This draws heavily on work first developed by Kelvin Lancaster [9] and Gary Becker, who assume that the household contains a production relationship that translates activities and goods into characteristics – and it is those characteristics that individuals actually value. This model has been applied many times in the past, to a number of different questions, and has led to the development of a sub-area of micro-economic analysis known as hedonic modeling. Becker's model in particular is useful for us, in that he assumes that one of the key inputs into the generation of the utility attributes is an individual's time, which must be allocated to each attribute (along with other inputs which represent market goods and services) subject to a global time constraint (e.g., number of hours in the day). An additional modification, which is made purely for expositional purposes, is that utility is a linear combination of levels of each domain.

2.2. Utility as Subjective Well-Being

Modern microeconomic analysis is largely built upon the notion that consumers make decisions by maximizing utility.[1] Given a small set of assumptions regarding the nature of individuals' preferences, one can demonstrate that a function can be derived which will represent a person's preferences by rank ordering all possible states of the world (usually expressed as different consumption bundles). This function forms the basis for most models of consumer behavior and is usually expressed as some variant of $U = u(\underline{\mathbf{x}})$ where $\underline{\mathbf{x}}$ is a vector of goods. While the general tendency was to view utility as an abstraction, economists have long estimated specific functional forms – for example, the McFadden Random Utility Model.[3]

The literature is enormous, and a meaningful overview is not within the scope of this section, but one aspect of empirical applications of utility theory is particularly relevant for this paper. This is the increasing interest economists are showing in re-integrating the concept of utility with the psychological construct of subjective well-being.[4-6] This literature on subjective well-being takes seriously the notion that in maximizing utility individuals seek to maximize a sense of happiness or life satisfaction, that this life satisfaction is quantifiable, and that interpersonal comparisons are possible (at least for broad policy analyses). In many ways, this literature represents a return to the conception of utility put forward by Jeremy Bentham and Francis Edgeworth.

In the model developed here, we draw specifically on the work by van Praag and colleagues who have proposed that overall utility – "general satisfaction" – can be modeled as a (linear) combination of the satisfaction that individuals receive in multiple "domain satisfactions." In their model, global satisfaction (GS) may be expressed as:

$$\textbf{(1)} \quad GS = GS(DS_1, DS_2, ..., DS_J; Z)$$

where DS_j represent individual domain satisfactions, and Z is a vector of explanatory variables, and where

$$\textbf{(2)} \quad DS_j = DS_j(X_j, Z)$$

and X_j is a vector of individual characteristics affecting the specific domains.[4] Domains include satisfaction with job, financial situation, housing, health, leisure, and so forth. One important extension of this work is that if income is included in the Z vector of explanatory factors for global satisfaction, then the parameter on income will reveal the individual's marginal valuation of income. The inverse of this parameter in a linear model is the dollar value of one unit of global satisfaction. Empirically, this method provides a direct means of assigning dollar values to overall life satisfaction (given the individual's characteristics) and will permit dollar values to be assigned to any changes in global satisfaction from, for example, changes in obesity. [6]

2.3. Adaptation

We will extend this notion of global satisfaction as the sum of domain satisfactions to incorporate adaptation. Previous work on estimating global satisfaction models has found evidence in favor of our adaptation modification. For example, Frijters (2004) notes that in his models respondents "tended to find the areas of their lives they are dissatisfied with less important." This is the sort of evidence one would expect to find if, as we will assume below, people are able to invest effort into adapting to circumstances and thus can react to negative outcomes either by changing the outcome (which is costly in terms of effort and money expended to purchase inputs to the outcome production process) or changing how much they care about the outcome (which is also costly in terms of effort expended to adjust expectations, peer groups, and the like).

There is widespread evidence of adaptation that is starting to make its way into the economics literature. In addition to the small effects of income on well-being (Easterlin 2001), there is evidence that the income an individual considers to be 'sufficient' is primarily determined by her current income (van Praag and Ferrer-I-Carbonell, 2004), and that adaptation appears to offset about two-thirds of the benefits of any increase in income (Frey and Stutzer, 2002). Using data from a 15-year study of over 24,000 Germans, Lucas et al (2003) show that, on average, people experience an increase in happiness in the years surrounding marriage but after the second year of marriage they appear to return to their baseline. Even in the case of widowhood, adaptation is close to complete after about eight years. It is worth noting that the same German data also highlight the point that adaptation is not found for all conditions. In the case of unemployment, for example, average life satisfaction falls from around 7.2 on a scale from 1-10 to 6.3 in the first year and is still only 6,.5 in the fourth year of unemployment unemployment (Clark et al, 2004).

It could be that some of the results suggesting adaptation are explained by response shift (Sprangers and Schwartz, 1999). Paraplegics, for example, might compare their happiness to other paraplegics, elevate their current ratings to reflect the contrast with the extreme despair immediately following the onset of disability, or adopt lower standards for the intensity of positive affect, all of which would lead to over interpretation of the degree of adaptation. However, whilst response shift makes intertemporal and interpersonal comparisons of self-reports problematic, it cannot explain all changes in preferences that take place. For example, there is strong evidence of adaptation even when physiological or behavioural measures are used, both of which should be less prone to response shift: Krupat (1974) found that that prior exposure to threat reduced galvanic skin conductance (a physiological measure of threat); and Dar et al (1995) found that war veterans with more severe past injuries could hold their finger in hot water for longer before classifying it as painful than veterans with less severe past injuries.

3. THEORETICAL MODEL

As discussed above, we will build upon the standard economic model of discounted utility maximization. Our approach will deviate from the standard model in three ways. We will assume: i) that individuals maximize utility across important life domains, such as health, work, leisure, social relationships etc.; ii) that utility is a linear combination of the levels of each domain, based upon the utility weight that that the individuals assigns to each domain; and iii) that the utility weights are themselves subject to modification by the individual, typically as an adaptive response to changed circumstances. For example, an individual may give less weight to the health domain – consider it to be less important to their overall utility – as their body mass increases, their cardiovascular health decreases or if they suffer some sudden shock such as a spinal cord injury. (Note that in principle this suggests that utility maximization may be better modeled as a dynamic process, where current utility depends on lagged values of some factors. However, for this first model we will ignore explicitly dynamic issues and explore a traditional static model, leaving a more complex dynamic model for future research.)

3.1. Defining Value Domains

The usual model assumes that agents possess a single preference relation, $\succeq$, which supports a one-dimensional ranking across all (pairwise) comparisons of commodity bundles $(x_i, x_j) \in X^N$, where X^N represents the commodity space, and x_i, x_j represent N-dimensional vectors of specific points in the overall commodity space X^N. However, for our model of adaptive utility, we assume that there exists a set of preference relations $\{\succeq_k\}$ where k=1 . . . K. Let this particular preference profile be defined as P ≡ ($\succeq_1, \succeq_2, \ldots, \succeq_K$), which represents the primitive preferences for a specific individual that maps the single bundle of commodities, x, onto rankings that are defined across the K life-satisfaction domains, which are designated by subscripts. Thus, rather than having one primitive, as in the standard model, we assume each individual has a set of K primitives, that correspond to “domains” of satisfaction (or well-being, or happiness, or utility).

Further, let each of the preference relations possess the following properties:

1. Completeness: $\forall$ pairs $(x_i, x_j) \in X^N$, it must be true that $x_i \succeq_k x_j$ or $x_j \succeq_k x_i$, or both;
2. Transitivity: $\forall$ $(x_i, x_j, x_l) \in X^N$, if $x_i \succeq_k x_j$ and $x_j \succeq_k x_l$, then $x_i \succeq_k x_l$;
3. Reflexivity: $\forall$ $x_i \in X^N$, $x_i \succeq_k x_i$;
4. Continuity: $\forall$ $x_i \subset X^N$, the sets $P(x_i \mid x_i \succeq_k x_0)$ and $I(x_i \mid x_0 \succeq_k x_i)$ are closed
5. Local Non-Satiation (strict monotonicity): $\forall$ pairs $(x_i, x_j) \in X^N$, $x_i \geq x_j \Rightarrow x_i \succeq_k x_j$.

Note, however, that these relations only obtain within a particular preference relation in the set of all preference relations simultaneously held by the agent. Thus, if $x_i \succeq_k x_j$ and $x_j \succeq_k x_l$ it does not follow that $x_i \succeq_m x_l$ when $m \neq k$. In other words, a given ranking of bundles that is supported by preferences in one domain will not generally apply to the other preference domains that an agent possesses. For example, a bundle that contains a house with larger rooms but a less attractive school district may be preferred in the housing domain but not preferred in any other domains.

With these restrictions on the nature of the set of preference relations held by the agent, it must be the case that we can rationalize each preference relation using a functional relationship v^k: $X^N \rightarrow \Re^+$ such that if $x_i \succeq_k x_j$ then $v^k(x_i) \geq v^k(x_j)$. These functions $v^k(x_i)$ $\forall$ k=1…K can be rationalized as value functions for each of the primitive utility domains, such that the single vector of commodities, x_i , generates K measures of satisfaction. Each of these value functions is assumed to display the usual properties of being (at least) twice differentiable, invariant to monotonic transformation (though in this case, we will need to restrict admissible transformations to be those that are applied equally to all 1, …, K domains simultaneously), and capable of representing the underlying preference orderings. What remains to identify the composite function $U(v^k(x_i))$ that represents the ultimate value function. For this, we must introduce a new primitive – a composition technology that will combine the domain value functions into a global utility function.

3.2. Defining Global Utility

With the set of K preference relations, P, defined, we assume that agents aggregate the supported value functionals, $v^k(x_i)$ into a composition function which defines the ultimate happiness. Thus, in this model, global utility is an aggregation of sub-"utility" functions. However, global utility must be constructed in a manner which preserves rationality. To assure this, we impose the following additional restrictions on the nature of revealed global preference behavior.

6. The global utility function is Paretian: $\forall$ U: P $\rightarrow \Re^+$ and for any alternative commodity bundles, x_i and x_j it must hold be true that $U(x_i) > U(x_j)$ if $x_i \succeq_k x_j$ $\forall$ k=1 … K.
7. The global utility function is generalized utilitarian in form: $\forall$ permissible U: P $\rightarrow \Re^+$, it is the case that U(x) = $\sum_k \omega_k (v^k(x))$, where $g_k(\ .\ )$ is increasing and quasi-concave.

These two additional restrictions on the nature of permissible global utility functions preserve the rationality of orderings. With generalized utilitarian functions (which are themselves Paretian), rankings of different commodity bundles will be transitive and complete. A linear composition is a special case of the generalized utilitarian form – defined as we have across quasi-concave sub-functionals. Given its notational advantage, we will explore the implications of our adaptive global utility model assuming a linear composition from this point forward. The results we obtain are found in a more general composition – though with the cost of substantial notational complexity. Thus, the model we explore below is defined as U(x) = $\sum_k \omega_k(\cdot) v^k(x)$.

A final modification to the standard model that we will introduce is to incorporate the possibility that the weights, $\omega_k(\ .\ )$, attached to particular domains of life may change as the circumstances associated with those domains change. So, for example, the importance an individual attaches to the sub-domain of health may change as a result of paraplegia. This can be seen in the wider context of the well established phenomenon of adaptation. Adaptation is a ubiquitous process for people, and can take many forms, including lowered expectations and substantive goal adjustment (Menzel at al, 2002). We make no pejorative claims about the adaptation process. Rather, this notion of adaptation fits well into the framework of utility maximization by assuming that an individual's adaptive processes affect the weights on each life domain that are used to construct the linear utility function.

We assume, therefore, that the utility weights attached to each life domain are themselves a function of individual choices – and subject to adjustment. Specifically, an individual must expend effort to maintain, or increase or decrease, how much they value a particular domain of utility. This is essentially the insight provided by Becker (1966) and Lancaster (1967). In order to convert commodities into the raw stuff of utility – characteristics, in the Becker/Lancaster language – an individual must apply some household technology. However, unlike Becker and Lancaster, we assume that the household technology is not free, but that opportunity costs must be incurred to utilize it. In the spirit of the literature on household production, we will assume that the opportunity cost of utilizing the household technology is the application of (non-market) "effort." In our model, then, we assume that investments of effort, e_k are allocated to maintaining / adjusting each of an individual's utility weights, $\omega^k(e_k)$, and that this effort must be allocated according to a maximum time / effort constraint:

$$(5)\quad E - L = \sum_{k=1}^{K} e_k$$

where E corresponds to the maximum amount of time/effort available and L corresponds to the time/effort devoted to the labor market to generate income. Note that labor, L, is simply another form of effort, e, but one which has a negative opportunity cost in the form of market wages – so that it will be convenient to distinguish it separately in our notation. Effort devoted to labor has a return equal to a market wage, m, and every other component of effort has an associated opportunity cost, π_k. Consistent with our definition of a generalized utilitarian global utility function, note that $\partial\omega^k(e_k)/\partial e_k \geq 0$, $\partial^2\omega^k(e_k)/\partial e_k^2 < 0$,and $\partial\omega^k(e_k)/\partial e_j = 0 \ \forall\ k \neq j$.

Thus, in a single period problem, a person will maximize the (static) utility

$$(6)\quad U_i = \sum_{k=1}^{K} \omega^k(e_k)\cdot v^k(x_{ik})\,,$$

subject to,

$$(7)\quad Income = I_0 + mE - \sum_{k=1}^{K}[m+\pi_k]e_k - \sum_{i=1}^{N} P_i \cdot x_i \quad = 0\,,$$

and where $x_i \in X^N$, m is the wage rate, π_k is the financial cost of effort devoted to non-labor activities (so that the full opportunity cost of allocating a unit of effort away from labor and toward attention to any utility sub-domain is = $m + \pi_{ik}$), and P_i correspond to market prices for the elements of x_i . Thus the agent's problem is to choose e_k, x_n, and λ in order to maximize

$$(8)\qquad L_i = \sum_{k=1}^{K} \omega^k(e_k)\cdot v^k(x_n) + \lambda\left\{ I_0 + mE - \sum_{k=1}^{K}[m+\pi_k]e_k - \sum_{i=1}^{N} P_i \cdot x_i \right\}$$

The F.O.C. for this problem are:

$$(9)\qquad L_{e_j} = \omega_j^j(e_j)v^j(x) - \lambda\left[m+\pi_j\right] = 0, \quad \forall \quad j = 1,\ldots,K$$

$$(10)\qquad L_{x_n} = \sum_{j=1}^{K} \omega^j(e_j)v_n^j(x) \quad - \quad \lambda P_n = 0, \quad \forall \quad n = 1,\ldots,N$$

$$(11)\qquad L_\lambda = I_0 + mE - \sum_{k=1}^{K}[m+\pi_k]e_k - \sum_{i=1}^{N} P_i \cdot x_i = 0$$

which has dimensionality of K + N + 1. The S.O.C. are the usual $\left|\overline{H}_2\right| > 0$, $\left|\overline{H}_3\right| < 0$, and

$$\left|\overline{H}_D\right| = \begin{vmatrix} \omega_{11}^1 v^1 & \cdots & 0 & \omega_1^1 v_1^1 & \cdots & \omega_1^1 v_N^1 & -[m+\pi_1] \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & \omega_{KK}^K v^K & \omega_K^K v_1^K & \cdots & \omega_K^K v_K^K & -[m+\pi_K] \\ \omega_1^1 v_1^1 & \cdots & \omega_K^1 v_1^1 & \sum_{j=1}^{K} \omega^j v_{11}^j & \cdots & \sum_{j=1}^{K} \omega^j v_{1N}^j & -P_1 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \omega_1^1 v_N^1 & \cdots & \omega_K^1 v_N^1 & \sum_{j=1}^{K} \omega^j v_{N1}^j & \cdots & \sum_{j=1}^{K} \omega^j v_{NN}^j & -P_N \\ -[m+\pi_1] & \cdots & -[m+\pi_K] & -P_1 & \cdots & -P_N & 0 \end{vmatrix} \begin{matrix} > \\ < \end{matrix}\ 0 \ as\, D\, is \begin{matrix} even \\ odd \end{matrix}$$

where D ≡ K + N + 1.

With the first and second order conditions, we can calculate the relevant comparative statics in order to discover how our adaptive global utility theory can contribute to understanding adaptive behaviors. First, solve implicitly for the optimum functionals for x(.), e(.) and λ(.), substitute these back into the F.O.C., and totally differentiate with respect to the parameters of interest. Taking the usual comparative static on price P_i for commodity x_n, we find as expected:

$$(12) \qquad \frac{\partial x_n}{\partial P_i} = \lambda \left|\overline{H}_D\right|^{-1} C_{K+i,K+n} - x_i \left|\overline{H}_D\right|^{-1} C_{D,K+n} = \lambda \left|\overline{H}_D\right|^{-1} C_{K+i,K+n} + x_i \frac{\partial x_n}{\partial I_0}$$

which is the standard Slutsky equation, where $C_{i,j}$ represents the cofactor for the i^{th} row and j^{th} column of the Jacobian of the system, and the first term is the negative substitution effect when n=i. Similarly, effort devoted to attending to the k^{th} domain in the global utility function also responds to changes in the price vector for commodities, as

$$(13) \qquad \frac{\partial e_j}{\partial P_i} = \lambda \left|\overline{H}_D\right|^{-1} C_{K+i,j} + x_i \frac{\partial e_j}{\partial I_0}.$$

However, unlike the standard model, the total effect of changing the price of a good is not confined to the simple comparative static on price. To see how, consider the full effect of the change in the own-price for good x_n. By totally differentiating the last FOC (11) with respect to P_n, solving for x_n and substituting back into (12), we get:

$$(14) \qquad \frac{\partial x_n}{\partial P_n} = \lambda \left|\overline{H}_D\right|^{-1} C_{K+n,K+n} - \frac{\partial x_n}{\partial I_0} \left\{ \sum_{j=1}^{K} \left[m + \pi_j\right] \frac{\partial e_j}{\partial P_n} + \sum_{\substack{i=1 \\ n \neq i}}^{N} P \frac{\partial x_i}{\partial P_n} \right\}$$

Again, while the first term is the always-negative own-price substitution effect, the magnitude of the income effect depends not only upon the responsiveness of the demand relationships of the other goods to the price of x_n, but it also depends upon the distribution of effort to each utility domain selected by the agent, and how those distributions are affected by the price of x_n. Using a similar approach, we can see that there is also an effect on the demand for x_n from changes to the opportunity cost of effort, e_j such that:

$$(15) \qquad \frac{\partial x_n}{\partial \pi_j} = \lambda \left| \overline{H}_D \right|^{-1} \sum_{j=1}^{K} C_{j,K+n} - \frac{\partial x_n}{\partial I_0} \left\{ \sum_{\substack{i=1 \\ i \neq j}}^{K} [m + \pi_i] \frac{\partial e_i}{\partial \pi_j} + \sum_{n=1}^{N} P \frac{\partial x_n}{\partial \pi_j} \right\}$$

Adaptation is further evidenced in this model if we optimize the last F.O.C. in (11) and then totally differentiate with respect to exogenous income, I_0, to arrive at:

$$(16) \qquad \sum_{j=1}^{K} [m + \pi_j] \frac{\partial e_j^*}{\partial I_0} + \sum_{n=1}^{N} P_n \frac{\partial x_n^*}{\partial I_0} = 1$$

This implies that unlike the standard model, an exogenous shock in income will in general *not* be completely allocated across the vector of commodities being consumed. Rather, some of the income increase will be allocated to the effort devoted toward maintaining the importance of the utility domains relative to one another. The comparative statics of effort with respect to income may be positive or negative – so that utility weights may actually fall as income increases. Consequently, the response of the actual level of utility to changes in income may be more muted that one would expect if only the commodity vector is taken into account.

4. DISCUSSION

The implications of this full model for research into individual decision making when adaptation is seen through the lens of an adaptive global utility function are significant. Generally, economists have considered elasticities of demand in terms of market prices – or at least on the direct opportunity costs of goods. This model suggests that the demand for a market good may be subject to a more complex set of determinants.

In particular, individuals' responses to own- or cross-good price changes will not only involve the direct substitution effect and traditional (quantity-weighted) income effect, but will also involve shifting the allocation of effort devoted to maintaining utility weights. This last effect can serve to attenuate (accentuate) the price elasticity of demand whenever the individual "utility production" technology is such that changing the importance of a utility domain in which the good plays a major role is relatively easy (more difficult). That is, if a price increase affects primarily a good that is very productive in $v^k(x)$, but for which the marginal productivity of e_k in $\omega^k(.)$ is low, then utility will be adjusted both by decreasing the amount of x_n consumed, and also by reducing e_k and shifting that effort to other domains, e_j, where the affected x plays a less central role.

Consider an example from the health field – obesity. Despite a range of efforts to encourage reductions in obesity, many people do not respond by lowering their body mass index (BMI). In the past, psychological researchers have focused on informational, motivational or dissonance oriented barriers to weight control or reduction. Economists have focused on the opportunity costs of weight control interventions. Our model suggests, however, that there are at least three classes of barriers to controlling weight and obesity – only one of which has been widely recognized and explored in the past.

First, individuals may find it difficult to control their BMI because the opportunity costs of doing so (in direct medical costs or foregone pleasure) are too high. This is the usual problem. Second, individuals may have difficulty in translating changes in BMI into changes in actual utility because of the limits in their individual production technology. It is possible that moderate changes in BMI may not lead to significant changes in each of the life domains that make up the ultimate utility function. Thus, the payoff from BMI reduction – even in the immediate term – may be too low to offset the opportunity costs of BMI reductions. Finally, individuals may ultimately fail to reduce BMI (even if it would lead to significant changes in the utility domains) because they may find it less costly to adapt to lower values of some utility domains. That is, the opportunity cost of reducing one or more of the utility weights on the domains where BMI plays a large role may be lower than the opportunity cost of actually changing BMI.

We welcome comments on any parts of the paper, of course, but would particularly like to get feedback on whether the model makes sense, whether we have accurately captured adaptation, what things we may have missed and what areas we might apply it to where empirical data could shed light on it explanatory ability.

Finally, we apologize for the lack of references but we were short of time and wanted to make sure some semblance of a paper was ready in time.