The Foundations of Utility for Quality of Service in Interface Development

Chris Johnson

Department of Computing Science, University of Glasgow, Scotland, G12 8QQ,
johnson@dcs.gla.ac.uk
WWW home page: http://www.dsc.gla.ac.uk/~johnson

Abstract. A number of research groups have exploited utility curves to model interaction with distributed systems. For example, they have been used to construct the models of subjective value that support "intelligent" advice giving systems. They have been integrated into ATM architectures to ensure that users' Quality of Service requirements are met by underlying network protocols. They have also been used to represent and reason about the risk aversion and risk preference that users exhibit when retrieving resources from remote servers over unreliable networks. However, much of this previous work has rested upon implicit assumptions about properties of the preference relation that underpins modern consumer theory. This paper examines the mathematical basis of the preference relation. The analysis helps to identify the implications that preference axioms have for the application of consumer theory to interface development.

1 Introduction

This paper builds on the notion that models of utility, or value, can be extended from microeconomics and consumer theory to model interaction with distributed systems. Utility can be thought of as a measure of 'desire' or the capacity of a good or service to satisfy a need [1]. The value of remote information can, therefore, be thought of in terms of its capacity to satisfy the demands that particular tasks impose upon users. A number of different approaches can be exploited to represent the utility that users associate with items of information. For instance, March and Simon's satisficing techniques identify the subjective 'desire' for a good or service by iteratively refining the constraining equations that characterise preferences between tasks and services [2]. Unfortunately, the difficulty of accurately identifying an individual's preferences has been a common theme of recent research in economics and consumer theory [3]. A number of research groups have, however, applied elements of utility models to support the development of interactive systems. For example, Horvitz and Rutledge have applied decision theoretic concepts to model users' actions under uncertainty [4]. McAuley's group at Microsoft Research in Cambridge are applying utility curves to model quality of service issues in ATM networks. Similarly, we have used consumer theory to model aspects of human computer interaction with distributed systems ranging from web-based education services [5] to safety-critical interfaces [1].

1.1 Assumptions About Utility

This previous work rests upon a number of implicit assumptions about the nature of utility. These assumptions have been directly inherited from the axioms that are used to characterise utility in the field of microeconomics [6]. No previous work has been conducted to examine whether or not these assumptions are valid within the specific context of human computer interaction. This paper, therefore, focuses in much greater detail on the mathematical foundations of consumer theory. In particular, we argue that the preference axioms of microeconomics make strong assumptions about patterns of interaction with mobile, distributed systems.

1.2 A Wireless LAN Case Study

Our analysis is illustrated by the problems of supporting interaction over a wireless local area network. This work is part of an on-going study into the use of personal digital assistants by Anaesthetists in the UK National Health Service. The particular architecture that we are using is based around Lucent's WaveLAN technology. This exploits a cellular architecture and spread-spectrum transmissions. It supports transparent migration between cells that are effective up to two hundred metres from a base station. This case study is appropriate because it typifies the sorts of distributed interactive systems in which users must sacrifice finite resources, both of time and battery power, to access the information that they need.

1.3 Structure of the Paper

This section has introduced the argument that is presented in this paper. Section 2 builds on this analysis by briefly presenting a number of utility curves that have been developed to characterise interaction with our wireless system. In particular, it is argued that appropriate feedback must be provided if relatively high levels of utility are to be sustained in the face of high network latencies [8]. Having illustrated the manner in which utility curves can be applied to represent and reason about interaction with mobile networks, Section 3 goes on to examine the preference relation that provides the formal basis of most utility curves. In particular, we examine the way in which mathematical properties of this relation that have been introduced from consumer theory may not be appropriate for models of human computer interaction. Section 4 builds on this analysis and uses a number of simple topologies to introduce the further axioms. These are necessary in order to move from the simple preference relation of Section 3 to the more familiar models of utility curves that have previously been used to model user preferences during human computer interaction. Later sections then argue that many of these axioms of 'rationality' and 'taste' cannot easily be applied to model interaction with distribued systems such as our WaveLAN case study. Section 5 presents the conclusions that can be drawn from this work and identifies areas for further research.

2 Using Models of Utility for Interface Development

Before exploring the formal model that supports consumer theory in microeconomics, it is first important to provide a brief overview of how this technique can be used to represent and reason about human computer interaction with distributed systems.

2.1 Indications of Cost in Mobile Networks

Figure 1 shows how utility curves can be used to analyse the delays that occur when many independent users share the same mobile infrastructure. Increases

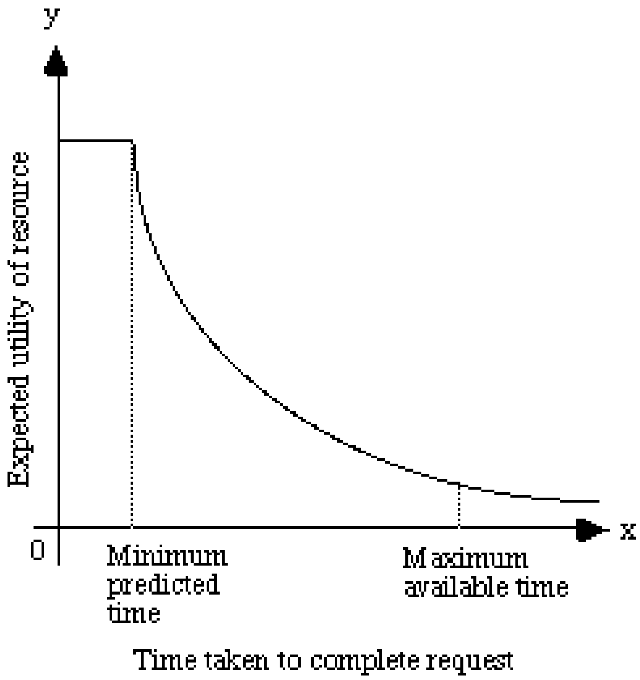

Fig. 1. Expected Utility in Distributed Systems

in network loading help to increase the costs, in terms of the users' time, that must be met in order to retrieve a remote resource. Initially, the user has some anticipated or predicted minimum for the amount of time that it will take to complete any transfer. Up to this point, the utility of any information retrieved will be relatively high. Documents received before the minimum predicted period can be seen as a bonus from the user's perspective. We show this as a plateau. However, it may be modelled as an extension of the "convex to the origin" curve if relatively small delays affect the utility that the user associates with the resource. After the minimum predicted time, the anticipated value of the information will decline. There are many reasons for this. The time taken to retrieve a document may exceed the total time that is available for any associated tasks. Eventually, there will come a point when users simply run out of time. At this stage, the marginal utility of any remote information will be very low because the user simply cannot afford to wait for the data to be retrieved [7].

It is possible to reduce the problems created by network delays by extending the period of high utility that is associated with the interval before the minimum expected retrieval period. However, the delay that is predicted by the user can be very different from the actual delay. This will be effected by transient conditions, such as the loading of any intermediate networks, and by local conditions associated with the users' position in relation to their nearest transceiver. Figure 2 presents a utility curve that shows how the minimum predicted time can be extended by a transformation along the X axis. In practical terms, this can be done by providing users with an indication of the capacity of their communications channels, by the size of the resource being transferred and also be the contents of their cache [5]. Ebling argues that this form of translucence is necessary if users are to maximise the allocation of their finite resources during interaction with mobile systems, such as our WaveLAN case study [8]. If the available network capacity is relatively low and the resource is relatively large then users will expect a considerabled delay before the information is available. Under these circumstances, the minimum expected time will be extended and high levels of utility will be associated with any request that completes before this interval has finished. All of this depends upon the user receiving appropriate feedback both about the size of the requested resources and about network latencies. One of the simplest means of providing this feedback using our mobile architecture is to support one of the commercial web browsers that provides information about the progress of retrieval requests. Unfortunately, multi-threaded Java implementation techniques easily defeats attempts to provide accurate assessments of download latencies in the general case. Further problems arise because it can be difficult to make a priori predictions about available bandwidth in mobile systems. For instance, communications 'shadows' and multipath transmissions can introduce delays while error correction routines call for the retransmission of corrupted or missing data.

2.2 Indications of Value

Instead of increasing the minimum expected download latency, it is possible to sustain high levels of subjective satisfaction in the face of wireless communica-

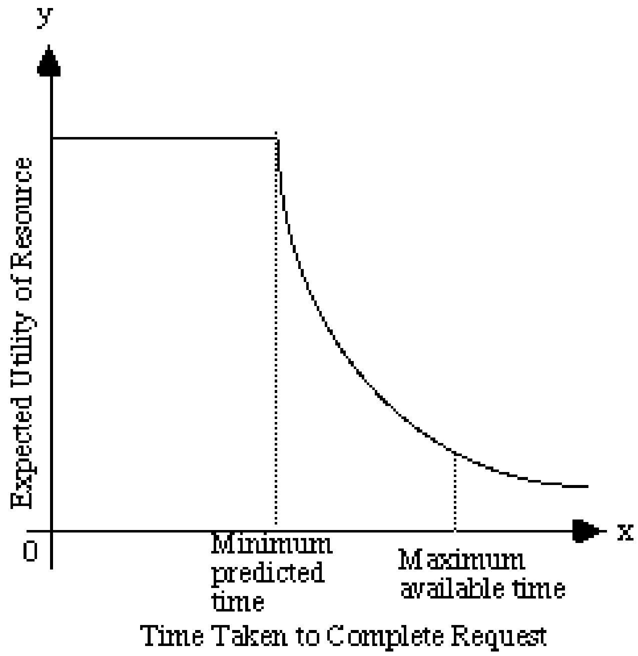

Fig. 2. Indications of Cost Extend Minimum Predictions

tions delays by providing users with information about the quality of a remote resource. An earlier paper explains how this technique can be applied to support the presentation of video material over distributed networks [9]. Significantly more requests are abandoned if hypertext tags are used to represent remote resources rather than thumb-nail images. Figure 3 presents a utility curve that can be used to explain these experimental findings. Initial assessments of subjective value, shown as a dotted line, can be transformed along the Y-axis. This indicates the rise in expectations that can be produced through effective interface design. Conversely, an interface which indicated relatively low production values would have the opposite effect of reducing the absolute level of the anticipated utility. This can be shown by a transformation of the utility curve down the Y axis.

This section has provided a brief overview of the ways in which utility curves can be extended from microeconomic consumer theory to model interaction with distributed systems. It is important not to underestimate the consequences of this analysis. Previous paragraphs have focused on the impact that subjective utility has for the presentation of information about download latencies and the production quality of remote resources. The introduction has also cited attempts to explicitly encode users' subjective utility as part of underlying network proto-

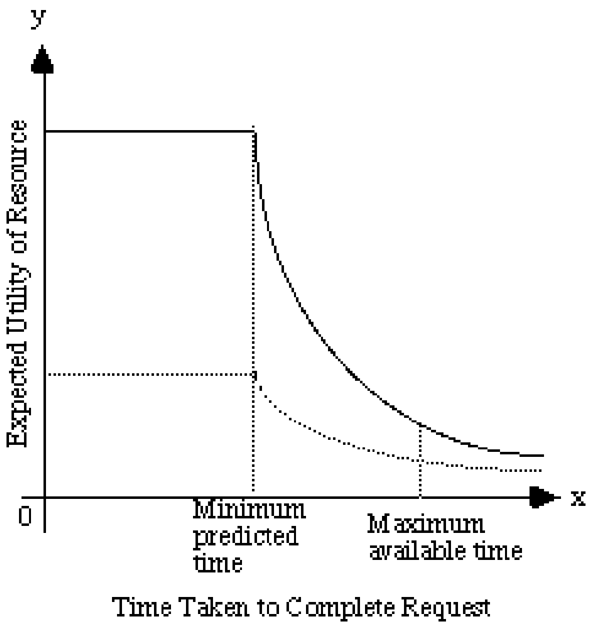

Fig. 3. Indications of Quality Raise Expected Utility

cols and in the user models that are being developed to drive the next generation of Microsoft's help systems [10].

3 The Preference Relation and Axioms of Rationality

The previous section applied utility functions to represent and reason about interaction with distributed systems. It did not, however, provide any mathematical basis for these functions. Nor did it provide any of the axioms that might distinguish 'valid' utility curves from those that characterise 'impossible' preferences [6]. This is important because without these axioms it is possible to inadvertently represent a number of paradoxical situations. For instance, a user might prefer to allocate finite network connectivity to access a web page rather than read their email and at the same time prefer to read their email rather than access the same web page. This form of 'deadly embrace' indicates a pathological form of inconsistency that cannot easily be explained using experimental observations about the way that real users choose to allocate battery life and radio connections. The following section, therefore, goes on to examine the theoretical foundations of consumer theory. These foundations are built on what are known as the rationality axioms of the preference relation. An analysis of these axioms helps to identifies important differences between new applications of utility to support interface development and its more 'traditional' uses in microeconomics.

Most work on consumer theory begins by defining a number of axioms that are assumed to hold over a consumption set X. A consumption set represents the set of all alternatives, or complete consumption plans, that the user can identify. For instance, a user may choose to download a particular item of information over their wireless connection. The user chooses to sacrifice some portion of their available battery power and their communications bandwidth in order to "consume" this resource. It is, therefore, possible to identify the remote item of information as a potential element of the user's consumption set. It is important to emphasise that it may not be possible for the user actually to consume particular elements of X. They may attempt to access information that cannot be obtained given their current resources. This often arises when the users' PDA has insufficient power to access the wireless LAN or if they move beyond the boundaries of the cellular architecture. In such circumstances, the request must be cached until the batteries are recharged or the user moves back into a cell.

It is possible to define a number of preference axioms that are assumed to hold over the elements of X. These axioms form the basic premises that are assumed to hold for most of the more complex utility models that have been exploited by previous work in this area. The preference axioms are intended to be a minimal set of relations that must hold in order to provide a consistent model of users' preferences within a consumption set. Consumer preferences are denoted by a binary relation over the elements of X. If $(x_1, x_2) \in \succeq$, or $x_1 \succeq x_2$, then "x_1 is at least as good as x_2". The $\succeq$ can, for example, be used to denote that a user prefers to allocate finite resources to retrieve one item of information, such as a web page, rather than another, for instance an email message. Although this relatively simple relationship only holds between two commodities, later sections will show how more complex topologies can be constructed from the initial formalisation. We will also show how microeconomic theory tends to assume a partial ordering over the elements of X. we can also impose further constraints on the $\succeq$ relation to avoid the paradoxes that were introduced in previous paragraphs.

3.1 Completeness

The first of the preference axioms ensures completeness over the elements in X. People must have the necessary knowledge to discriminate between the different commodities that they are presented with. They must be able to state that for any two distinct consumption plans, x_1 and x_2, that x_1 is at least as good as x_2 or x_2 is at least as good as x_1.

Axiom 1 *For any $x_1 \neq x_2 \in X$, either $x_1 \succeq x_2$ or $x_2 \succeq x_1$*

As mentioned, the reason for examining the formal underpinnings of consumer theory is to uncover the implications that this model has for its application to human computer interaction. The Completeness Axiom makes a number of assumptions about the users of distributed applications. It is intended to ensure that users must be able to distinguish between the different strategies or plans that they can exploit during interaction. For example, if a user's task involves two remote items of information but they cannot determine whether they have sufficient power to retrieve either resource then it can be difficult to express a preference between them. In order to make an informed selection they will need more information about probable power consumption. It will be difficult to sustain the preference relation, $\succeq$, under these circumstances.

HCI Implication 1 *The Completeness Axiom makes an unrealistic assumption that the user will be able to distinguish between the different strategies or plans that they can exploit during interaction.*

Experience in the development of mobile applications shows that without significant training many users cannot accurately distinguish between the options that are available to them. Even relatively frequent users have only the crudest notion of caching or of cellular transmission protocols. As a result, a high number of relatively 'poor' exchanges are made in which battery power is extinguished by attempts to transfer documents that could have been deferred until conventional network connectivity was restored.

3.2 Reflexivity

The second preference axioms also expresses an important principle for the rational decision maker. In particular, it ensures that any good or service is at least as good as itself. Without such an axiom, it is possible to envisage a considerable number of paradoxes. The reflexivity axiom is, therefore, formalised as follows:

Axiom 2 *For all $x \in X$, $x \succeq x$*

In terms of human computer interaction, this ensures that the value associated with a remote resource is at least as good as the value of that remote resource. However, even this relatively benign axiom has certain problematic features when applied to interaction with distributed systems.

HCI Implication 2 *The Reflexivity Axiom states that the value of any good or service is at least as good as itself. This implies that each data source must be modelled as a distinct commodity even if the information that they contain is identical.*

The key point here is that there may be several different ways in which to obtain the same remote resource. Ideally, if information is cached then any request can be handled locally. However, if the same resource is not cached then the costs of accessing that resource are very different. Most people would prefer to access the locally cached version than the remote resource with its associated network connection and battery consumption. This leads us to the conclusion that we must be very careful about what we model as different commodities within any application of utility theory to distributed systems. The Reflexivity Axiom teaches us that local and remote versions of the same information are different resources. They must not be confused when applying models from microeconomics to represent and reason about human computer interaction.

3.3 Transitivity

The transitivity axiom helps to build up the partial ordering that reflect many consumer attitudes to the commodities that are available to them:

Axiom 3 *For any* $x_1, x_2, x_3 \in X$ *if* $x_1 \succeq x_2$ *and* $x_2 \succeq x_3$ *then* $x_1 \succeq x_3$

As before, it is possible to identify ways in which the transitivity axiom helps to model certain aspects of interaction with distributed systems. For instance, if a mobile user preferred to access their email account rather than download a web page and they prefer to download the web page rather than remotely print out the list of patients allocated to an operating theatre then it is also safe to conclude that they would prefer to cache their email than print the list. It is important to note, however, that this is the most controversial of the preference axioms. The previous example illustrates the point that these initial axioms characterise the notion of the "rational consumer". They describe an ideal that can be very different from the insights provided by empirical studies.

HCI Implication 3 *The Transitivity Axiom makes an unrealistic assumption that users act as "rational" consumers in a technical environment that they may not fully understand and which contains numerous distractions and delays.*

A number of experimental studies have shown that individuals do not always express preferences that are consistent with the transitivity axiom [3]. This is especially true when users hold only 'mild' preferences between individual consumption patterns. Inserting a small delay or distraction before asking people about such orderings can lead users to express preferences that violate the transitivity axiom. Unfortunately, such delays and distractions are typical of the working environments of most mobile users. This has profound implications. Several mobile systems have been developed to automatically 'optimise' users' caching and information transfer tasks [8, 11]. Such approaches, typically, assume a degree of rationality that is a key feature of the transitivity axiom but which is not always exhibited by observational studies of mobile human computer interaction [12].

4 Preference Topologies and the Axioms of Taste

Previous sections have introduced the axioms that characterise the preference relation in microeconomic consumer theory. These axioms have been implicitly adopted by previous work that has sought to apply time dependent models of utility to user interface design [1]. In order to illustrate this connection, it is necessary to show how the previous axioms can be used both to constrain the $\succeq$ relation and also to develop topologies that form the more familiar utility curves shown in Figures 1, 2 and 3. This section also introduces further axioms that characterise and constrain these utility curves.

Definition 1. *The $\succeq$ ('preference') relation on the consumption set X is assumed to satisfy Axioms 1, 2 and 3.*

We can use the Completeness, Reflexivity and Transitivity axioms to define a strict preference relation. This excludes the possibility that any two commodities might be equally preferred. It can be used to express total orderings amongst the elements of X.

Definition 2. $x_1 \succ x_2$ *('strict preference') if and only if* $x_1 \succeq x_2$ *and* $x_1 \not\preceq x_2$

This is important because there are often situations in which users are confident in expressing marked preferences between certain remote resources. For instance, we have developed a system that enables its users to express a strong preference that they receive messages from particular colleagues, such as their boss, before any other messages. The ability to express such preferences has important consequences for the effective deployment of finite bandwidth and battery power.

In contrast to strict preference, it is also possible to construct an indifference relation. This captures the idea that there will be some commodities for which the user cannot express a preference. Recall, however, that this need not violate the completeness axiom which defines a partial rather than a total ordering. The indifference relation captures situations in which users do not care which of several remote servers might be chosen to address a particular information request.

Definition 3. $x_1 \sim x_2$ *('indifference') if and only if* $x_1 \succeq x_2$ *and* $x_1 \succeq x_2$

The previous definitions enable us to develop a topology with respect to any arbitrary commodity, x_0, that the user might select. In particular, we can define subsets of X such that the elements of that set are at least as good as that point, $\{x|x \in X, x \succeq x_0\}$. We can define a "no better than" subset, $\{x|x \in X, x_0 \succeq x\}$. Similarly, a "worse than" subset, $\{x|x \in X, x_0 \succ x\}$. "Preferred to" is defined as follows: $\{x|x \in X, x \succ x_0\}$. Finally, the indifference subset is given as: $\{x|x \in X, x \sim x_0\}$. It is, difficult to visualise the components of this topology if we think in terms of remote information resources that can be accessed over a wireless network. Figure 4, therefore, shows X to be composed of 2D vectors of infinite real numbers $X = R^n_+$ [6]. As can be seen, there is an area above the point $x_0 = (x^0_1, x^0_2)$ such that commodities in that region will be preferable to those represented by x_0. Conversely, the area below x_0 represents commodities that are no better than this point. Finally, there is an indifferent area to the left of x_0 that contains commodities which are valued as highly as the goods that are represented by this point. The key point here is that the first three preference axioms which help us to define preference relations also permit a number of anomalous topologies, such as that indicated in Figure 4. For example, there

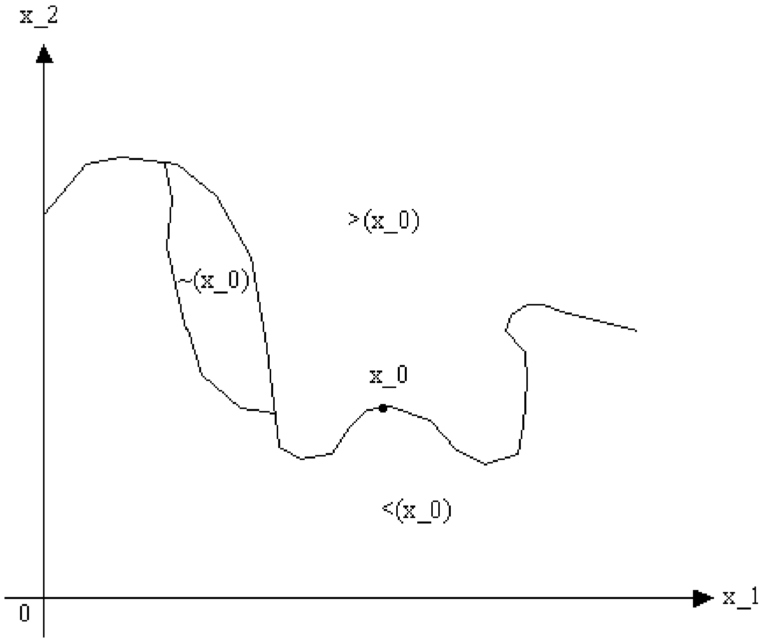

Fig. 4. Example Preference Topology Defined By Axioms 1, 2 and 3

is a small region immediately to the left of x_0 where the user would actually prefer to have reductions in both x_1 and x_2. If we consider that x_1 and x_2 might represent a trade-off between using power to download incoming email messages and using power to transmit outgoing messages then this leads to a paradox where the user would prefer both to transmit and receive less messages. This situation can only be explained by enriching our notion of X to multidimensional vectors that consider wider resources implications; for example the user might choose not to access their mail because they want to conserve power for other applications. Alternatively, such paradoxes can be avoided by introducing further axioms that, like the others already presented, are often implicit in the concepts of utility that have been used in previous studies of human computer interaction. These axioms further constrain the preference relation to the point at which the resulting topologies will more closely resemble the conventional utility curves, such as those shown in Figures 1, 2 and 3.

4.1 Continuity

The continuity axiom ensures the closure of $\succeq (x)$, which is read as "at least as good as" and the $\preceq (x)$ relation, which is read as "no better than".

Axiom 4 *For all $x \in R^n_+$, the "at least as good as" and "no better than" sets are closed in R^n_+.*

Informally, a set in R^n is open if it does not contain any of the points on its boundary. It is closed if it does contain all points on its boundary. Stating that both "at least as good as" and "no better than" are closed, ensures that their complements are open. In other words, the boundary points between "at least as good as", $\succeq (x)$, and "worse than", $\prec (x)$, are entirely contained within "at least as good as". Conversely, the boundary points between "no better than", $\preceq (x)$, and "preferred to", $\succ (x)$, are entirely contained in "no better than".

HCI Implication 4 *The Continuity Axiom ensures topological niceity and is neutral with respect to Human Computer Interaction.*

Our analysis has failed to find any immediate implications of the Continuity Axiom for Human Computer Interaction. However, it does have important theoretical implications for the development of consumer theory. In particular, it helps to ensure that sudden reversals of preference do not occur [6].

4.2 Strict Monotonicity

The remaining axioms that can be applied to the preference relation are intended to reflect particular assumptions about consumer, or user, tastes. In most consumer theory there is an assumption that more is always better than less. This has become formalised by the axiom of strict monotonicity. It can be expressed as an assumption that if one bundle of commodity x contains at least as much of every commodity in y, then the first bundle is at least as good as the second.

Axiom 5 *For all $x_0, x_1 \in R^n_+$ if $x_0 \geq x_1$ then $x_0 \succeq x1$ while if $X_0 \gg x_1$ then $X_0 \succ x_1$.*

This axiom together with the continuity axiom also help to avoid the paradox in Section 4. If the elements of the consumption set were represented by vectors that modelled the amount of incoming and outgoing email then users would always prefer situations in which there were more incoming and outgoing messages. However, this interpretation again illustrates some of the concerns that arise when investigating the mathematical foundations of consumer theory.

HCI Implication 5 *The Axiom of Strict Monotonicity fails to characterise certain aspects of mobile interaction in which more communication can rationally be seen to be a worse situation than one in which connectivity is reduced.*

This observation might, at first, be seen to be paradoxical. However, a number of authors have written about the potential disruption that is caused by bringing mobile telephones and other networked computing devices into complex working environments [13]. Other authors have written about the social consequences of users being forced to reveal their geographical location in wireless LANs as they move from cell to cell [14]. These social and contextual factors can lead rational users to deliberately disconnect from a network. In other words, they may prefer less connectivity rather than more and hence may reject the Axiom of Strict Monotonicity.

Previous paragraphs have argued that the Axiom of Strict Monotonicity need not be regarded in the same way as the previous 4 axioms because its role in consumer theory is to model user 'tastes' towards particular commodities. It should, therefore, not be surprising that it is inappropriate for certain aspects of human computer interaction with distributed systems. With this in mind, Figure 5 illustrates the impact that Axiom 5 has upon upon the topology that was introduced in Figure 4. As can be seen, continuity and then mononicity first

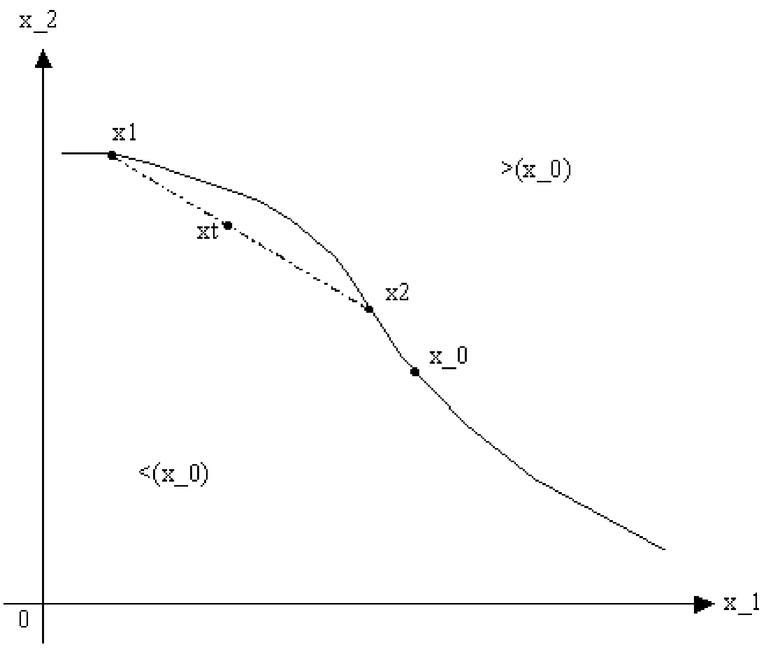

Fig. 5. Example Preference Topology Satisfying Axioms 1,2,3,4 and 5

bound and then reduce the indifference region that was shown in the previous example. Axiom 5 ensures that that the indifference sets in R^2_+ cannot bend upwards or contain positively sloped segments and that all preferred sets are strictly above indifference sets.

4.3 Strict Convexity

The final axiom that is, typically, introduced to model preferences between commodities is known as Strict Convexity. It is intended to ensure that the resulting topologies assume the same "convex to the origin" shape that can be seen in the utility curves of Figures 1, 2 and 3. The impact of this axiom can be seen by looking at the indifference points $x1$ and $x2$ in Figure 5. These are indifferent because they lie on the same boundary between $\succ(x)$ and $\prec(x)$ that $x0$ occupies [6]. From this we can conclude that $x1 \succeq x2$, however, convex combinations of these points, such as xt will lie within $\prec(x0)$ and hence violate Axiom 6.

Axiom 6 *If $x_1 \neq x_0$ and $x_1 \succeq x_0$ then $t\ x_1 + (1-t)x_0 \succ x_0$ for all $t \in (0,1)$*

In terms of microeconomic theory, Axiom 6 leads to the principle of diminishing marginal rate of substitution in consumption. One way of thinking about the implications of Strict Convexity is that it reflects a predisposition or bias on the part of the user to prefer "mixed" bundles of consumables rather than the extremes that appear at either end of the curve in Figure 5. As a result, this axiom models users who take a balanced view of trading finite time or battery power for remote information. If they had little battery power left then they would be less likely to exhaust it with further network interaction than if they had a plentiful supply.

HCI Implication 6 *The Axiom of Strict Convexity reflects a "balanced" approach to resource allocation or substitution. As one of the preference axioms of 'taste', however, it need not be appropriate as a model for all forms of human computer interaction.*

Again, however, it is possible to identify exceptions to this axiom. For instance, we have seen instances in which users will trade the last of their battery life to access a remote resource. This represents a novel form of what Reason has referred to as encysting within other domains of human computer interaction [15]. The user becomes so preoccupied with particular sub-tasks that they lose their overall awareness of the overall success or failure of their interaction.

5 Conclusion

A number of research groups have exploited utility curves to model interaction with distributed systems. They have been used to develop the models of subjective value that support "intelligent" advice giving systems [10]. They have also been used to represent and reason about the risk aversion and risk preference that users exhibit when retrieving resources from remote servers over unreliable networks [5]. Finally, they have been integrated into ATM architectures to ensure that users' requirements are met by the underlying network protocols. However, much of this previous work has rested upon implicit assumptions about properties of the preference relation that underpins modern consumer theory. This paper has examined the mathematical basis of the preference relation. The analysis helps to identify the implications that preference axioms have for the application of consumer theory to interface development. In particular, we have argued that:

1. The Completeness Axiom makes an unrealistic assumption that the user will be able to distinguish between the different strategies or plans that they can exploit during interaction.
2. The Reflexivity Axiom states that the value of any good or service is at least as good as itself. This implies that each data source must be modelled as a distinct commodity even if the information that they contain is identical.
3. The Transitivity Axiom cannot easily be sustained because empirical studies have shown that preference paradoxes often arise when users do not fully understand their technical environment or when users are faced with numerous distractions and delays.
4. The Continuity Axiom ensures topological niceity and is neutral with respect to Human Computer Interaction.
5. The Axiom of Strict Monotonicity fails to characterise certain aspects of mobile interaction in which more communication can rationally be seen to be a worse situation than one in which connectivity is reduced.
6. The Axiom of Strict Convexity reflects a "balanced" approach to resource allocation or substitution. As one of the preference axioms of 'taste', however, it need not be appropriate as a model for all forms of human computer interaction.

There is, however, one important caveat. We have not argued for a one-to-one relationship between the preference relation, $\succeq$, and utility curves, such as those shown in Figures 1, 2 and 3. The preference relation defines an ordering over the elements of the consumption set, X. Utility functions simply represent those preferences by mapping preference rankings onto a simple numeric scale. More preferred commodities are assigned higher values while less preferred commodities are assigned lower numbers. Indifferent commodities are assigned equivalent values. Having raised this caveat, it is important to state that the utility function reflects the complete, reflexive, transitive, continuous and strictly monotonic properties of the preference relation. Hence our analysis applies both to the preference relation and to the utility functions that reflect that relation.

Utility functions are increasingly being used as a tool to model human computer interaction. Our analysis has shown that there are a number of theoretical issues that remain to be addressed before these techniques can accurately capture the diverse patterns of interaction that are observed in many mobile applications. In particular, our initial experience with developing interfaces for wireless LANs has shown that the axioms of rationality often do not apply to users who are struggling to come to terms with new styles of interaction. Similarly, the axioms of taste only go part of the way towards constructing valid models of the ways in which users choose to expend finite battery power and communications bandwidth in mobile environments.

References

1. C.W. Johnson Decision Theory And Safety-Critical Interfaces In K. Nordby and P. H. Helmersen and D. Gilmore and S.A. Arnesen (eds.) Human Computer Interaction - Interact '95, Chapman and Hall, London, United Kingdom, 127-132, 1995.

2. J.G. March and H.A. Simon, Organisations, Wiley, New York, United States of America, 1958.
3. C. Puppe Distorted Probabilities And Choice Under Risk, Springer Verlag, Lecture Notes In Economics And Mathematical Systems, No 363, Berlin, Germany, 1991.
4. E. Horvitz and G. Rutledge, Time-Dependent Utility and Action Under Uncertainty, Proc. 7th Conf. on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Mateo, CA, 1991, pp. 151-158.
5. C.W. Johnson (1995), Time and the Web: Representing and Reasoning about Temporal Properties of Interaction with Distributed Systems. In M. Kirby, A. Dix and J. Finlay (eds.), People and Computers X, 39-50, Cambridge University Press, Cambridge, United Kingdom.
6. G.A. Jehle and P.J. Reny, Microeconomic Theory Addison Wesley, Reading, MA, 1998.
7. P. O'Donnel and S. Draper.: How Machine Delays Change User Strategies. Adjunct Proceedings to HCI'95, British Computer Society, London. 1995.
8. M.R. Ebling, Translucent Cache Management for Mobile Computing. Doctoral Dissertation, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, March 1998.
9. C.W. Johnson The Ten Golden Rules for Providing Video over the Web. In C. Forsythe, E. Grose and J. Ratner (eds.) Human Factors and the Web, Lawrence Erlbaum, New York, United States of America, 1998.
10. E. Horvitz and T. Paek, A Computational Architecture for Conversation. In J. Kay (ed.), Proceedings of User Modelling 1999. Springer Verlag, 1999.
11. N. Davies, G. Blair, K. Cheverst and A. Friday, Supporting Adaptive Services in a Heterogeneous Mobile Environment. In L.-F. Cabrera and M. Satyanarayanan (eds.), Proc. Workshop on Mobile Computing Systems and Applications (MCSA) Santa Cruz, CA, U.S., IEEE Computer Society Press, Pages 153-157. December 1994.
12. K. Vaananen-Vainio-Mattila and S. Ruuska, User Needs for Mobile Communication Devices: Requirements Gathering and Analysis through Contextual Inquiry In C.W. Johnson (ed.), Proc. of the First Workshop on HCI and Mobile Devices. Department of Computing Science, Glasgow University. Technical Report (GAAG-98-1), 1998.
13. P. Johnson, Exploiting Context in HCI Design for Mobile Systems, In C.W. Johnson (ed.), Proc. of the First Workshop on HCI and Mobile Devices. Department of Computing Science, Glasgow University. Technical Report (GAAG-98-1), 1998.
14. C.W. Johnson, The Impact of Time and Place on the Operation of Mobile Computing Devices, In H. Thimbleby, B. O'Conaill and P. Thomas (eds.), People and Computers XII: Proceedings of HCI'97, Springer Verlag, London, United Kingdom, 175-190, 1997 .
15. J. Reason, Human Error, Cambridge University Press, Cambridge, UK, 1980.