

2008

Homeward Bound: Ecological Design of Domestic Information Systems

David Wastell

Nottingham Business School, UK, david.wastell@nottingham.ac.uk

Follow this and additional works at: <http://aisel.aisnet.org/ecis2008>

Recommended Citation

Wastell, David, "Homeward Bound: Ecological Design of Domestic Information Systems" (2008). *ECIS 2008 Proceedings*. 169.
<http://aisel.aisnet.org/ecis2008/169>

This material is brought to you by the European Conference on Information Systems (ECIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ECIS 2008 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

HOMeward BOUND: ECOLOGICAL DESIGN OF DOMESTIC INFORMATION SYSTEMS

Wastell, David, Nottingham University Business School, Jubilee Campus, Wollaton Road, Nottingham, UK, david.wastell@nottingham.ac.uk

Abstract

IT artefacts are steadily permeating everyday life, just as they have colonized the business domain. Although research in our field has largely addressed the workplace, researchers are beginning to take an interest in the home environment too. Here, we address the domestic realm, focusing on the design of complex, interactive information systems. As such, our work sits in the design science version rather than behavioural science paradigm of IS research. We argue that the home is in many ways a more challenging environment for the designer than the workplace, making good design of critical importance. Regrettably, the opposite would appear to be the norm. Two experiments are reported, both concerned with the design of the user interface for domestic heating systems. Of note is our use of a medium-fidelity laboratory simulation or “microworld” in this work. Two main substantive findings resulted. First that ecologically designed feedback, embodying a strong mapping between task goals and system status, produced superior task performance. Second, that predictive decision aids provided clear benefits over other forms of user support, such as advisory systems. General implications for the design of domestic information systems are discussed, followed by reflections on the nature of design work in IS, and on the design science project itself. It is concluded that the microworld approach has considerable potential for developing IS design theory. The methodological challenges of design research are highlighted, especially the presence of additional validity threats posed by the need to construct artefacts in order to evaluate theory. It is argued that design theory is necessarily complex, modal and uncertain, and that design science (like design itself) should be prosecuted in an open, heuristic spirit, drawing more on the proven methods of “good design” (e.g. prototyping, user participation) in terms of its own praxis.

Keywords: design science, cognitive ergonomics, ecological design, domestic heating system, feedback, operator performance, goal setting theory.

1 INTRODUCTION

This paper is about the design of IS artefacts in the home. Although ICT pervades more and more of contemporary life, research in our field remains pre-occupied with the workplace (Brown and Venkatesh, 2005; Venkatesh and Brown, 2001). In neighbouring disciplines, such as human-computer-interaction and cognitive engineering, there is smattering of design research addressing ICT in the domestic sphere. The few studies in our discipline typically address contextual factors influencing technology adoption (e.g. Brown and Venkatesh, 2005; Vijayasathy, 2004). Whilst the importance of design factors (ease of use, usefulness etc.) is acknowledged (e.g. Brown and Venkatesh, 2005) such “adoption studies” intrinsically take the artefact as the object of passive study, not something to be actively and directly shaped. Beliefs and attitudes towards technology, and how these bear on intentions “to use or not to use”, form the quintessential field of interest. Design is an element in this nexus but only one component in a welter of other causal and contextual variables (social norms, demographics and so on). The typical theoretical framing of these factor studies draws either directly on the diffusionism of Rogers (1995), or on derivative positions such as the ubiquitous Technology Acceptance Model (Davis et al, 1989).

The classic adoption study fall squarely within what Hevner calls the behavioural science approach to IS research, in contrast to the design science paradigm (Hevner et al., 2004; March and Smith, 1995). The same broad dichotomy re-appears in Gregor’s typology of IS theory (Gregor, 2006), which contrasts theory for explanation and prediction with theory for design and action. What would seem to underlie the distinction is the degree to which the design of an artefact is a matter of immediate concern to the researcher, or whether the relationship is passive and deferred. Although a well-established tradition with a lineage reaching back over many years (e.g. Markus et al., 2002; March and Smith, 1995; Walls et al., 1992), IS design research remains a minority genre. By reporting an example, we hope to redress the balance and strengthen the presence of design research within the IS repertoire. By emphasizing the importance of design, we make common cause with cognitive ergonomists (notably Norman, 1998) who have inveighed against the proliferation of technical devices that are increasingly complex, opaque and frustrating to use.

1.1 Designing for the home

The design of information systems for the effective management of domestic central heating systems (CHS) provides our substantive focus. We shall describe and discuss two contiguous cycles of a design project funded by the German Research Council on ecological design. The overall goal of our research programme is to explore the potential for energy savings through improving the design of domestic artefacts. CHS has by far the largest environmental impact of all technologies in the home and is therefore a priority concern. Typically, the management of the CHS is mediated by a crude information system, in the form of a periodic energy bill which provides cumulative information regarding energy consumption and costs. The potential for computerizing this information system, for exploring different modes of information presentation and decision support, is of obvious relevance, not only in terms of domestic economics but the broader ecological agenda.

Designing for the home poses some unique challenges, compared to the design of work-based systems. This brings us to the second sense of “bound” in the paper’s title: the group of users is characterized by a high level of heterogeneity, without the possibility of selection according to technical competence; moreover, no formal training can be given, and users in effect set themselves their own goals and tasks, with no performance supervision, standards or systematic feedback. In general, these contextual differences demonstrate the great importance of careful design since the potential for influencing behaviour is much more limited in the home than in a work context. In our investigation, we shall focus on the range of options for displaying information and aiding decision-making in order to achieve the optimal balance between comfort levels and energy expenditure.

Following Kroemer et al. (2001), four main display categories are distinguished: status displays (indicating the current system state), historical displays (information about past trends), predictive displays (projected information on future trends), and instructional displays (providing operational guidance). Research on the relative merits of different forms of display is extensive and widely scattered, clustering under various headings: HCI, cognitive ergonomics, cognitive engineering etc. We have not the space here to provide more than a very brief overview of the main relevant themes. As a general comment, it must be said that design absolutes are hard to find amongst the plethora of sometimes inconsistent or contradictory findings in this diverse literature.

Regarding status displays, some such feedback is essential for the operation of any device, but the choice of information to display is problematic where there are multiple interacting variables (Bennett et al., 2005). It might be thought that portraying historical trends is always advantageous, and such displays are indeed in widespread use. Yet the little empirical research that has been done has thrown up some negative findings (e.g. Palmer et al., 1980; Bennett et al., 2005). A less equivocal picture emerges for predictive displays, with benefits reported in many domains ranging from medicine to aviation (Wickens et al., 2000). Such tools are especially useful for managing lagged non-linear systems where the anticipation of future evolution is inherently complex (Wickens et al., 2000). Instructional displays provide qualitative feedback guiding operator action; experts systems provide a much-hyped example of the genre. Although such systems have the potential to aid decision-making in constrained, highly structured settings (Wickens and Hollands, 2000), there is often considerable resistance to their use in professional domains. Medicine is one such area, where clear evidence in terms of improved decision outcomes is notably wanting (Sintchenko and Coeira, 2003). This resistance recalls the disinclination noted in our field to the use of decision support systems such as executive information systems (Elam and Leidner, 1995). Unless perfect reliability of the decision aid can be assumed, there is evidence that simple status displays may actually be preferable to instructional displays which can give misleading advice (Sarter and Schroeder, 2001).

Two design experiments are reported here. In the first, the potential benefits of historical displays are evaluated, together with enhanced “ecological” feedback. The second experiment assessed the relative advantages of predictive and instructional displays. The findings of the two studies will be presented and discussed, before moving on to a more general set of reflections on the design science programme, motivated by recent debate over its nature and agenda (Hevner et al., 2005). As an instance of design science in action, the work offers the opportunity to reflect critically on the practice of design, and indeed design science itself. There is specific interest in the investigative approach used here, which has some innovative aspects. It involves an experimental methodology incorporating a dynamic, computer-based laboratory simulation. Such “microworlds” are invaluable for investigating complex, often inaccessible, settings where direct observation of key phenomena is problematic (Brehmer and Dörner, 1993). Designing for the home is methodologically challenging. Whilst a clear understanding of user requirements is essential, the designer is not in the advantaged position of his organizational counterpart who can directly co-opt users into the design process. Alternative, less immediate methods of accruing design knowledge must be sought. The microworld is an attempt to insert the real world into the “arc of design” at a formative stage, and it will be of interest to observe how effective this method is.

2 RESEARCH METHODOLOGY

2.1 Experiments as heuristic devices: introductory reflections

Whereas experimental research is normally associated with the rigorous version of positivism, the present work was carried out in a more heuristic spirit. Whilst empirical evaluation is an indispensable part of design science, the strong form of directional hypothesis testing would seem at odds with the inventive, creative spirit of design (Boland and Collopy, 2004). It would also seem to represent an

unrealistically optimistic view of the certainties of knowledge in the realm of the artificial, a world of contingent truths rather than necessary laws (as we glimpsed in the preceding section). Given the scientific uncertainties of the terrain, strong *a priori* hypotheses were eschewed. Beyond the broad expectation that more user support will, in general, lead to better performance we ventured forward with an open pragmatic mind, to appraise empirically what design concepts work and what do not, to attempt to understand some of the underlying contingencies and to abduct some tempered generalizations.

2.2 The CHESSE microworld

As noted, the experimental work was carried using a PC-based simulation of a generic CHS, dubbed CHESSE (Central HEating System Simulation). The design of CHESSE itself is of methodological relevance in relation to our general design science concerns. The version for experiment 1 was largely fashioned by the experimenters themselves (i.e. with no user involvement!) supplemented by expert feedback from engineering colleagues with interests in “green design”. Version 2 did involve an element of user participation, in that subjects from experiment 1 were debriefed about their experience and their comments led to several useful improvements in system usability. As with any microworld, verisimilitude was only carried as far as the pragmatics of the experiment required: time-scales were greatly accelerated, energy units were arbitrary, system dynamics were highly simplified, and of course the psychological mimicking of comfort levels could only be mediated in the crudest of ways. Within these limitations, CHESSE was designed to provide a realistic user experience, in terms of the abstract nature of the operator’s task, the visual appearance of the interface, the underlying dynamic properties of the heating system (heat losses, lags etc.). The aim was to produce a generic model that was convincing enough to generate meaningful user engagement and glean some relevant design knowledge regarding tool support and task performance.

CHESSE can be configured to represent a range of types of accommodation and heating arrangements. Specific temperature and weather profiles can also be created by the experimenter and stored prior to experimental sessions. CHESSE also creates a results file in which all key performance parameters (energy consumption, comfort levels etc.) are written separately for each “day” the simulation runs. All settings made and sources of information sampled by the operator are also logged in the results file. A small “one person” apartment with three rooms was used for both experiments. The operator’s main task was to define a heating profile for each of the rooms (i.e. sitting room, bedroom, and kitchen) according to a target temperature profile. There is a “set up” screen for each room, and the heating profile can be typed in via a dialogue box specifying the start, end and thermostatic temperature level for one or more “heating blocks”; alternatively the mouse may be used, drawing the block directly on a temperature by time-of-day graphic. Once the user is satisfied with the room settings s/he has made, the main simulation screen is selected (Figure 1). Clicking the “run simulation” button, fast-forwards the simulation for a complete day before pausing again; this takes around 30 seconds unless the operator decides to interrupt and intervene. Half-hourly status information on temperatures, cumulative energy use and comfort levels is provided on this screen, as well as a graphic showing diurnal trends for individual rooms.

Satisfactory comfort was defined as the attainment of a room temperature within 1 degree of the target. If the room temperature falls more than 1 degree below target but within 3 degrees, “mild discomfort” is indicated. A more severe discrepancy is signalled by a “serious discomfort” indicator. Figure 1 shows the situation for one participant at 23:54, when the day is nearly at a close. Only the bedroom has a target value. Clearly the heating has been switched off slightly prematurely, reflected in the disparity between target and actual temperatures marked by the appearance of the “mild discomfort” warning. A number of sources of management information were potentially available in both experiments at any point, depending on the experimental condition: a weather forecast provided a prediction of daily peak and low temperatures over the next four days. There was also a monthly report summarizing daily energy consumption and comfort levels for the house as a whole. Detailed

daily reports were also potentially available, showing the hourly temperature profile for each room (akin to figure 1) for any previous day, together with a summary of energy consumption and comfort.

3 EXPERIMENT 1

The first experiment compared the following three levels of operator support. The standard mode (STAND) was intended to represent the normal situation in most households: cumulative information only was provided in the form of an overall indication of total energy used at the current point (equivalent to reading the gas or electricity meter). Participants were also informed of how effective they had been in heating the house, i.e. the average comfort level achieved thus far that day. This information was provided in aggregate form on the monthly report, i.e. without any daily breakdown. In the historical mode (HIST), participants were given full access to detailed historical information for each day, via the monthly and daily report screens. They were able to inspect the temperature and heating profiles for any room on any selected day, together with the total energy and comfort achieved for that day.

In the ecological (ECOL) mode, participants were provided with enhanced feedback in the form of a “waste estimator”, which used a crude algorithm to estimate how much energy had been unnecessarily consumed by switching the heating on too soon, or switching it off too late. This indicator was intended to overcome obvious weaknesses in the basic status information available in the STAND and HIST modes. Although the two relevant status variables were presented (energy and comfort), it is not readily clear how they trade off against one another, nor how they are impacted by operator behavior. The waste indicator provides a simple example of feedback that is “ecological” (Rasmussen et al., 1994) in the sense that it enables the operator “to directly perceive the state of affairs in the environment... the deep structure of the work domain” and provides some guidance for adaptive action (i.e. reviewing switching on/off times).

3.1 Participants, design and procedure

45 participants (30.0% female) were recruited from the Darmstadt University student population. They were aged 19 to 38 years (mean = 24.3 yrs) and received a payment of 25 for their participation. A one-factorial between-subjects experimental design was employed, with operator support as the primary independent variable, varied at three levels, as above. The experiment took place in a laboratory setting, and participants were tested in groups of 3 or 4 working on individual PCs, separated by screens. The testing session comprised two main phases. In the first, they were introduced to the CHESS software and initial training was given. In the second “treatment” phase, participants completed two task scenarios in their assigned feedback condition, each scenario involving the operation of the heating system for a whole month. Sessions on average lasted around 3 hours. Reflecting the home context, participants were not assigned specific goals; they were instructed “to do their best” to achieve the comfort levels prescribed by the set target temperatures whilst minimizing energy consumption.

User performance was measured on three dimensions: energy consumption, comfort (as defined above), and energy efficiency. Energy consumption was measured as the amount of energy (expressed as the average cost per day) consumed during task completion. The energy waste index was used to assess efficiency, although it was only fed back to participants in the ECOL condition. Patterns of information sampling and intervention were also assessed throughout. The frequency of inspection of the monthly and daily reports, the room settings and the weather forecast was measured. System control activity was assessed by measuring the frequency of changes to the heating profiles.

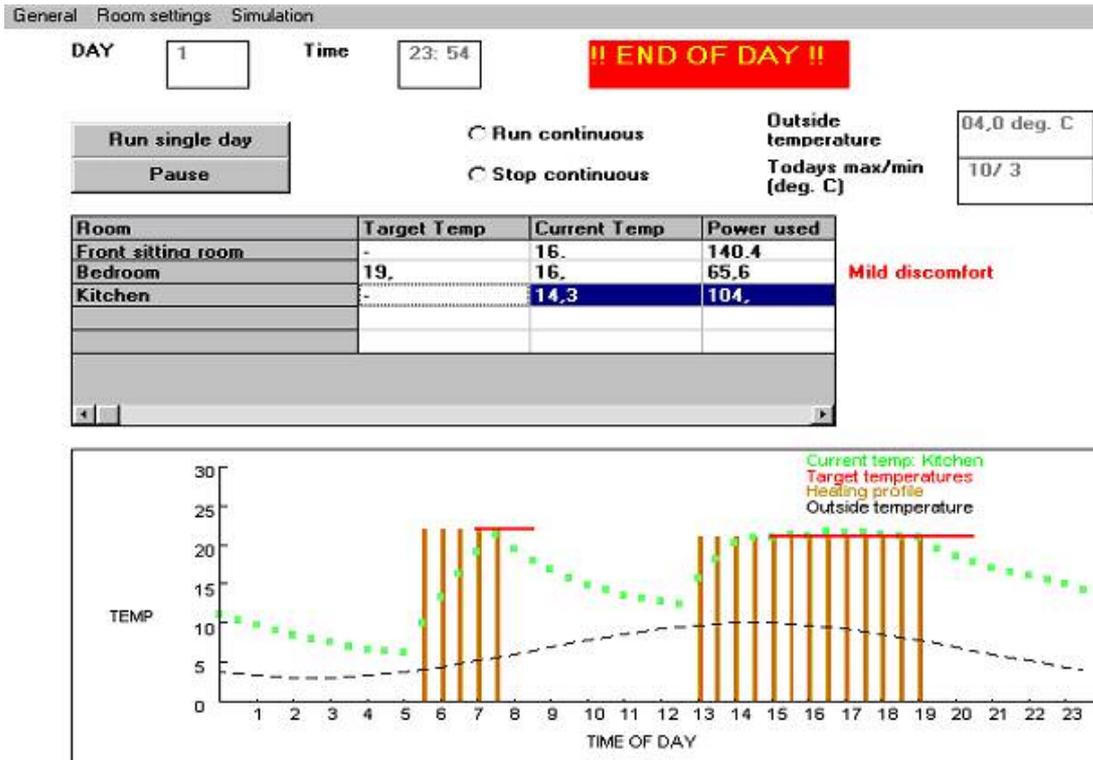


Figure 1. The main simulation display is shown at the end of a simulated day.

3.2 Results

Given the multiplicity of dependent variables (9 in all, including the 5 measures of user interaction) a two factor MANOVA was carried out as the first step in the formal statistical analysis, with operator support (feedback mode) and day-of-the-month as the two independent variables. Only the second treatment month was analyzed, as participants were at their most experienced and proficient at this point. Wilks lambda was significant for both support ($F=16.9$, $p<0.000$) and day ($F=27.7$, $p<0.000$). Table 1 shows the impact of operator support, in terms of group average performance levels across the month. Univariate F ratios were significant for all parameters. Tukey post-hoc tests indicated that participants achieved significantly higher levels of comfort in both the HIST and ECOL conditions, and that their energy use and waste levels were significantly reduced in the ECOL condition. All other differences were not significant. It is apparent that whilst the HIST group achieved significantly higher comfort levels than the standard group, they used more energy to realize this. Their waste index was also the same as the standard group. Participants in the ECOL condition, on the other hand, achieved improved comfort, but with no significant increase in energy use, in fact a small reduction was noted. Their waste index was less than half the other 2 groups!

Feedback Mode	Energy cost per day	Comfort (max 100%)	Est. waste
Ecological (ECOL)	5.4	93.8	5.3
Historical (HIST)	5.7	92.6	13.4
Standard (STAN)	5.6	88.5	12.8
F ratio	11.8	26.7	53.2

Table 1. Effects of feedback level on key performance indicators (Experiment 1).

Table 2 summarizes patterns of user interaction with the simulation. First we note that there is a low level of interaction with the system overall, generally less than 1 interaction per day per option. Looking at the monthly trend, it was found that most interaction occurred at the outset of the month, as would be expected, thereafter rapidly falling away and achieving steady state at around day 6. Room screens, for instance, were accessed 2.7 times on average on day 1 (to set/adjust heating profiles); by day 8, this had dropped off to less than once per day (0.7). Tukey HSD post-hoc tests revealed two themes. First, that for all forms of interaction, user activity was at its most prolific in the ECOL condition. Second, that for temperature forecasts and changes to heating profiles, there was consistently higher interaction in the STAND than the HIST condition, whereas for monthly reports and room screens, there were no differences between these two feedback conditions.

Feedback mode	STAN	HIST	ECOL	F ratio
Sampling of monthly report (No./day)	0.28	0.25	0.40	11.6
Sampling of historical daily reports	N/A	0.16	0.26	12.5
Accessing room screens (No./day)	0.90	0.74	1.31	14.8
Changes to profile settings (No./day)	0.82	0.60	1.07	14.9
Sampling of weather forecast	0.87	0.55	1.31	23.5

Table 2: *Information sampling and system intervention as a function of feedback mode (Experiment 1). All F ratios are based on $df = 2, 1318$ (pooled error term).*

4 EXPERIMENT 2

The second experiment differed from the first in two substantive respects. First, in the provision of more sophisticated support tools: a predictive display and an instructional adviser, both optional features. The second difference was the posing of a more demanding heating management task, including an external temperature profile that fluctuated significantly more markedly.

Three levels of operator support were provided. The first condition (HIST2) was a direct replication of the HIST condition of experiment one. The other two conditions involved additional support tools for setting room heating profiles. The first of these (PRED) provided a predictive display showing the estimated temperature profile for that room, given the weather forecast and the heating blocks defined by the operator. Building on the results of experiment 1, an indicator of energy-efficiency was also provided. This gave a comparison of the predicted performance against the optimal that could be achieved for a given set of target and forecast external temperatures. In the third condition (INST), an instructional display was also available, providing explicit feedback on efficiency and comfort as well as advice on how to improve performance (e.g. to switch on sooner and/or to set a lower target temperatures).

4.1 Participants, design and procedure

Forty-five new participants (48.0% female) took part in the study, again from the Darmstadt University student population. Their ages ranged from 19 to 47 years (mean = 23.7 yrs). They were paid 15 each for their participation, and general procedures were the same as experiment 1. Sessions lasted approximately two hours. After a general introduction to the experiment, participants were introduced to the CHESS software, and received practical instruction in its use. They then completed one test session involving the operation of the heating system for 30 days, which took approximately 60 min; this involved a complex sequence of daily scenarios modelled on typical lifestyle motifs: a

normal working day, an extended working day, a weekend day including a late party, a day at home. Each scenario had a specific heating profile to be implemented.

4.2 Results

Again, MANOVA was performed with experimental condition (3 groups) and day-of-the-month as the independent variables. This found that both factors were significant, yielding a Wilks lambda of 34.5 ($p < 0.000$) for the main effect of user support. Table 3 shows that whereas the participants with access to the predictive display were able to cope adequately with the more exacting challenge of this experiment, the performance of participants with only historical information was significantly impaired in terms of their ability to attain satisfactory comfort levels. Univariate F tests were significant for all three parameters in table 3, with Tukey HSD post-hoc comparisons indicating that the participants in the HIST2 mode, though more parsimonious in their energy use, were less effective in attaining comfort levels and showed a generally lower level of energy efficiency. No differences in performance were found between the PRED and INST groups for any parameter.

User support mode	Daily energy cost	Comfort	Energy-Efficiency
History (HIST2)	6.0	74.4	12.4
Predictive (PRED)	6.3	90.3	14.3
Instructional (INST)	6.3	91.6	14.5
F ratio	13.8	86.5	97.6

Table 3. Average energy consumption, comfort and energy efficiency levels for the second experiment. Note that energy-efficiency is simply comfort divided by daily energy cost.

Table 4 summarizes patterns of operator interaction. The need for more intensive monitoring and intervention is borne out by these figures, which are notably elevated compared to their counterparts in experiment 1 (they show similar intensity to the levels evidenced at the outset of the month in the earlier experiment). The general pattern across conditions is striking. There are two main themes. A more active profile of interaction is consistently seen for the HIST2 condition; this is especially remarkable for the temperature forecast and the historical reports which are hardly accessed at all in the PRED and INST conditions. The second theme is the very similar profile of interaction for the two conditions where the predictive display is available (PRED and INST). Tukey post hoc tests revealed this general pattern to be statistically valid for all parameters except the monthly reports, which were sampled more often in the PRED condition.

	HIST2	PRED	INST	F ratio
Sampling of monthly report (no./day)	0.34	0.45	0.13	43.2
Sampling of historical daily reports	2.8	0.7	0.3	57.7
Accessing room screens (no./day)	4.6	3.3	3.1	15.1
Changes to profile settings (no./day)	3.3	2.6	2.8	28.4
Sampling of weather forecast (no./day)	1.8	0.3	0.1	117.2

Table 4. Information sampling and system intervention for experiment 2. All F ratios are based on $df = 2, 1318$ (pooled error term).

Regarding the use of the advanced support tools, there was striking evidence of a strong preference for the predictive display. This was accessed on average 10.3 times per day by participants in the INST condition, and 11.5 times in the PRED group (this difference was not significant, $t=0.18$). In contrast, the energy adviser was rather scantily consulted, with the average rate of use being less than once per experimental day (mean = 0.81). To triangulate these behavioural measures, participants were asked to rate their subjective opinions of the various information displays on a seven point scale, running from 1 (low utility) to 7 (very useful). The HIST2 group gave a significantly higher assessment of the value of the historical display (6.1) than either of the two other groups (4.2 and 3.7 on average for PRED and INST respectively, $F = 4.9$, $P < 0.01$). This group also gave a notably higher rating (4.3) to the weather forecast than either of the other conditions, who gave an average rating of 2.5 ($F=12.3$, $P < .001$). Both groups with the access to the predictive display rated it very highly (6.3 and 6.4 on average), markedly higher than the general evaluation of the instructional display (5.2 on average).

5 DISCUSSION

In contrast to the mainstream “behavioural” approach to research on the adoption of IT/IS, the present work takes a more militant, though complementary, tack. The aim is to influence directly the take-up of technology by improving its design, centring attention on the delivery of relevant functionality embodied in a supportive user interface. We also set foot in an application domain hitherto not much explored in our discipline, namely the home environment. We have noted that the home is, in many ways, a more challenging environment for the designer than the workplace, where supervision and training can help to optimize the utilization of IT. Without such support, it becomes of paramount importance that technology is well-designed if it is to be used effectively. But as others have lamented, all too often this is not the case; bad design would seem to be the norm for everyday objects (Norman, 1998). Our discussion will be organized under two broad headings, following the twin aims pre-figured in the Introduction. In this penultimate section, we will reflect directly on the specific knowledge furnished regarding IS design for the home. In the final section, will stand back and reflect more generally on design work and design science, based on our experiences across here.

Our main substantive findings can be expressed in the following two generalizations. First, that historical displays, whilst replete with detailed information, do not in themselves assist users to manage the heating system more effectively, whereas the provision of feedback directly indicating energy-efficiency strongly motivated more interaction, which in turn enhanced performance. Secondly, that decision aids in the form of predictive displays are well-liked by users, extensively used and engender enhanced outcomes, whereas instructional support was relatively disregarded. Concerning the first finding, it is notable that the HIST group in experiment one actually intervened less often than the standard group, despite the radical improvement in their management information. By contrast, the addition of a single database field (the waste index) in the ECOL condition evoked a *gestalt* change in operator behaviour: all information sources were sampled more frequently, including the weather forecast (strongly indicative of a more pro-active orientation). Increased levels of system intervention were also observed. These changes in engagement directly translated into more energy-efficient system management, with comfort levels and energy consumption being jointly optimized. How may this dramatic transformation be explained?

For an explanation, we first return to the concept of ecological information systems design (EISD) briefly alluded to above (Rasmussen et al, 1994). The essence of EISD is the imperative to support human judgment in the management of complex problems through the clear representation of task objectives together with “measuring functions” enabling the direct perception of critical environmental variables and relationships (Rasmussen et al, 1994). Although the *a priori* motivation for the waste indicator had been merely to provide additional feedback, it was realized retrospectively that its real potency lay elsewhere, i.e. that it provided an implicit performance goal. The powerful influence of goals over performance is well known. Goal-setting theory (Locke and Latham, 2002) argues that specific performance goals are the primary regulator of task performance. A clear mapping between

goals and performance feedback is thus vital, though often lacking, a deficit which Norman (1998) memorably dubs the “Gulf of Evaluation”. Setting specific external goals is, however, unrealistic for the domestic environment where goals are voluntary and self-imposed. Although not designed as a goal-setting mechanism (and no explicit waste goals were set), the waste indicator nonetheless carried a potent implicit goal in its moral labelling (all waste is bad and should be reduced to zero). Not only does it set a goal but it directly conveys performance feedback against that goal, combining information from two countervailing parameters (comfort and cost) into a single unequivocal index. No mental computation is required to gauge performance; the error signal is directly given. Combining goal and feedback, it provides a simple example of ecological design neatly bridging the Gulf of Evaluation. We thus regard the superior performance of the ECOL group as a cameo validation of the philosophy of EISD, though we re-emphasize that the success of the waste-indicator was largely serendipitous; its teleological properties were only appreciated *ex post facto*.

Looking back at the first experiment, we can read the pattern of results as a general commentary on the contingencies of design and the limitations of design agency (Richardson, 1993). Designers sometimes get it right, but sometimes they do not. Both forms of enhanced feedback were expected to enhance system management, but only one afforded genuine benefit. The same generic pattern comes through in the second study: one aid worked whereas another failed to deliver the expected gains. The predictive display was well-liked by users, extensively used and enhanced operator performance. This result is consistent with the broad swathe of research which demonstrates the clear benefits of predictive displays in managing systems that are complex, lagged and dynamic (Wickens et al., 2000).

The findings concerning decision aids broadly concur with the main themes of the literature briefly reviewed in the Introduction. Whilst arguing for the clear potential advantages of predictive displays, this literature is markedly more equivocal on the benefits of other forms of decision support (Wickens et al., 2000). Users in our simulated home environment are apparently just as disinclined to utilize advisory systems as are professional users in real work settings (Sintchenko and Coeira, 2003). Although consistent with a literature that is (to say the least) equivocal on the benefits of prescriptive aids, definitive conclusions should be drawn carefully. The failure could, of course, simply reflect a poorly designed implementation rather than a fundamental flaw in the concept of instructional support. Nonetheless, we may safely infer that the design of such support is more challenging than predictive aids, and that there is a greater range of contingencies to be addressed. We may conclude that designers of complex artefacts for the home, and indeed for other environments which involve the management of dynamic processes, should, in the first instance, devote their efforts towards the provision of effective predictive displays, and exercise caution before investing design effort on other forms of support.

6 CODA: REFLECTIONS ON DESIGN AND DESIGN SCIENCE

As well as an experiment for design, the CHES studies constitute an experiment on the process of designing itself. Let us begin these final reflections by considering the insights into design work (and its potential dysfunctions) revealed by the studies, before considering more strategic implications for design science. First, we note that design is not a top-down, linear process moving deductively from a body of *a priori* knowledge (Boland and Collopy, 2004). Its open, emergent character is clearly manifest in the first experiment, where the full properties of an important design feature (the waste indicator) came to light adventitiously, and the relevance of a significant body of theory also only became fully apparent once the experimental work had been completed (i.e. Goal-setting theory). The tendency in design work to incorporate novel features that please the intellectual curiosity of the designer, rather than the pragmatic needs of users, has been dubbed “creeping featurism” by Norman (1998). This phenomenon is present in both experiments, and would seem to be an endemic feature of design! Significant effort was devoted in the first experiment to the development of the historical database, which provided no real benefit. In the second experiment, the expert system was relatively under-utilized. We have noted above that there was little user involvement in our design process, and

the salutary thought occurs that had we consulted our users more, it is possible that such redundant features would have been abandoned, or perhaps refined to be more useful.

An important subset of design theory concerns the efficacy and aptness of design methods and tools, and the performance of the microworld was a matter of obvious interest. The general value of such microworlds in the context of IS research has been argued by Wastell (1997). By inserting something of the complexity of the “real world” into the controllable space of the laboratory, it is possible to engineer and explore realistic and important scenarios that would be difficult to inspect *in vivo*. It also becomes feasible to evaluate formally the influence of a range of environmental, task and design variables. Although artificial, the medium fidelity realism of the microworld manifestly has the dramaturgical power to engage users and to elicit realistic behaviour. Certainly our users took their work seriously; this is evident from the pattern of their interaction with the system. The fact that outcomes ran against our broad expectations in itself provides a cogent argument for the use of such laboratory studies. Of course, the simulation is not the real thing, and external validity is inevitably problematic, as for any laboratory experiment. Such limitations are certainly recognized here. Nonetheless, we believe (albeit with due circumspection) that some valid and useful design knowledge has been generated by the microworld, and that the trends seen here have some generalizability. Design science differs from other realms of science (March and Smith, 1995; Walls et al. 1992) in that its knowledge-base must be judged against utilitarian criteria as well as predictive and explanatory criteria. This utility test (Venable, 2006) has been passed; we have confirmed the heuristic potential of microworld methodology for generating realistic behavioural data, testing ideas and developing design theory.

Design science differs from everyday design work in its aspiration to produce theory, not to build specific working systems. The artefacts constructed here are simply the means to this loftier end, i.e. vehicles for testing ideas and hypotheses in the endeavour to construct a body of generalizable knowledge regarding the design of a certain class of artefact. Methodologically, we would certainly affirm design science to be challenging. The experimental psychologist deals with one primary source of validity threats in her theory-building labours, those arising from the quality of the experimental design. She has the luxury of working with highly simplified experimental models (decocting real-world phenomena to some putative canonical form) and the theoretical knowledge sought is relatively simple and well-circumscribed.

In contrast, empirical realism is vital in design science, enabling knowledge to generalize across a wide range of contexts of use, tempered by relevant singularities. We should therefore expect to see the accumulation of a relatively unstable body of contingent knowledge punctuated by exceptions and caveats (rather than necessary, timeless truths) whose codification will present an unending challenge! In this disorderly world, induction and abduction represent more relevant modes of reasoning than strict deduction, as researchers try to make sense of the complex and inconsistent findings furnished by even relatively simple experiments such as the present ones. Doing design science necessarily involves design work, regarding both the design of the “experimentation” (field or laboratory) and of the artefact itself. Like design, we have argued that it is best prosecuted heuristically and opportunistically, but perhaps this argument should be pushed further. There is much that reflects the *habitus* of conventional behavioural science in the present work. The experimenters were very much in charge, and the empirics had a largely linear trajectory with little adaptation to user feedback. Critical elements of good design practice were lacking, in particular prototyping and user participation. We will end on the chastening thought that, had we worked more collaboratively and iteratively with our users (using prototypes, for instance, to explore design ideas) we may well have produced not only a better artefact but more robust theory too. Reflecting this, users in our future work will be more fully engaged as partners and co-producers of design knowledge, rather than passive guinea pigs!

References

- Bennett, K.B., Payne, M. and Walyers, B. An evaluation of a "Time Tunnel" display format for the presentation of temporal information. *Human factors* (47:2), 2005, 342-359.
- Boland, R. and Collopy, F. (2004). *Managing as designing*. Stanford University Press.
- Brandon, G. and Lewis, A. Reducing household energy consumption: a qualitative and quantitative field study. *Journal of Environmental Psychology* (19), 1999, 75-85.
- Brehmer B. and Dörner, D. Experiments with computer-simulated microworlds: Escaping both the narrow straits of the laboratory and the deep blue sea of the field study. *Computers in Human Behavior* (9), 1993, 171-184.
- Brown, S.A. and Venkatesh, V. Model of adoption of technology in households: a baseline model test and extension incorporating household lifecycle. *MIS Quarterly* (29:3), 2005, 399-426.
- Davis, F., D., Bagozzi, R., and Warshaw, P., (1989), 'User acceptance of technology: a comparison of two theoretical models', *Management Science*, 35, pp. 982-1003.
- Gregor, S.. The Nature of theory in information systems. *MIS Quarterly* (30:3), 611-642, 2006.
- Hevner, A.R., March, S.T., Park, J. and Ram, S. Design science in information systems research. *MIS Quarterly* (28:1), 75-105, 2004.
- Kroemer, K., Kroemer, H. and Kroemer-Elbert, K. *Ergonomics: how to design for ease and efficiency*. Prentice-hall, New Jersey, 2001.
- Locke, E.A. and Latham, G.P. Building a practically useful theory of goal setting and task motivation: a 35 year odyssey. *American Psychologist* (57), 2002, 705-717.
- March, S.T. and Smith, G. Design and natural science research on information technology. *Decision support systems* (15), 1995, 251-266.
- Markus, L., Majchrzak, A. and Gasser, L. A design theory for systems that support emergent knowledge processes. *MIS Quarterly* (26), 2002, 179-212.
- Norman, D.A. *The design of everyday things*. MIT Press, London, 1988.
- Palmer, E.A., S. Jago, D.L. Baty and S. O'Connor, Perception of Horizontal Aircraft Separation on a Cockpit Display of Traffic Information. *Human Factors Journal* (22:5), 1980, 605-620.
- Petroski, H. *Success through failure: the paradox of design*. Princeton University Press, 2006.
- Rasmussen, J. Pejtersen, A.M. and Goodstein, L.P. *Cognitive systems engineering*. Wiley, New York, 1994.
- Richardson, A. The death of the designer. *Design Issues* (9), 1993, 34-43.
- Rogers, E. M., (1995), *Diffusion of innovations*, The Free Press, New York.
- Sarter, N.B. and Schroeder, B. Supporting decision-making and action selection under time pressure and uncertainty: the case of in-flight icing. *Human Factors* (43:4), 2001, 573-583.
- Simon, H. *The sciences of the artificial*. MIT Press: Cambridge, 1969.
- Sintchenko, V. and Coeira, E. Which clinical decisions benefit from automation: a task complexity approach. *International Journal of Medical Informatics* (70), 2003, 309-316.
- Venable, J.R. The role of theory and theorizing in design science research. In, *Proceedings of the 1st International Conference on Design Science Research in Information Systems and Technology*. Claremont Graduate University, California, 2006.
- Venkatesh, V. and Brown, S.A longitudinal investigation of personal computers in homes: adoption determinants and emerging challenges. *MIS Quarterly* (25:1), 2001, 71-102.
- Vijayasarathy, L.R. Predicting consumer intentions to use on-line shopping: the case for an augmented technology acceptance model, *Information and Management*, 2004, 747-762.
- Walls, J.G., Widmeyer, G.R. and el Sawy, O.A. Building an information systems design theory for vigilant EIS. *Information Systems Research* (3:1), 1992, 36-59.
- Wastell, D.G. Human-machine dynamics in complex information systems: the "microworld" paradigm as a heuristic tool for developing theory & exploring design issues. *Information Systems Journal* (6), 1997, 245-260