Modelling of individual domestic occupancy and energy demand behaviours using existing datasets and probabilistic modelling methods.

Graeme Flett*, Nick Kelly

Energy Systems Research Unit (ESRU), Department of Mechanical and Aerospace Engineering, University of Strathclyde, Glasgow, UK

Abstract

With growing trends towards smaller-scale, low carbon energy systems such as microgrids and later generation district heating, coupled with the increased use of building simulation in bottom-up stock modelling, there is increasing interest in being able to automatically generate multiple, high-resolution profiles of both occupancy and occupancy-dependent household demand. This has resulted in the emergence of a range of tools capable of producing time-varying occupancy and demand profiles that capture many of the characteristics evident in real-world data. However, as a result of limited data availability, these tools have typically been calibrated using composite data from multiple individuals or households. This results in the production of profiles which whilst statistically faithful to the average characteristics of the underpinning data, fail to capture the variance seen from household-to-household in the real world.

This paper attempts to address this shortcoming, developing and testing a new method within an existing model structure for the production of occupancy and linked demand data using a probabilistic model that is representative of the overall population from which it is derived, but which has an improved ability to generate specific outputs that better match the behaviour of specific households. The new approach utilises previously established occupancy and demand modelling methods, along with the composite population-based calibration data used within the models. However, the temporal predictive basis of the models has been manipulated to account for individual behaviours whilst retaining the overall statistical characteristics of the source data. This was achieved by adapting the Markov chain timing basis for previously developed models and factoring the probability values. A linked electrical demand model has also been adapted by manipulation of the relative timing of use of different appliances and demands to account for differences between individual and average behaviours. The described approach has the benefit of not requiring any additional calibration data, which for both occupancy and energy demand is often scarce. The predictions of the improved model and previous version are compared to real occupancy and demand data, indicating that the alterations enable significantly more diverse profiles to be generated, whilst still being representative of the supporting data.

Keywords: demand, occupancy modelling, disaggregated, building simulation, probabilistic model

1. Introduction

A strong relationship between building occupancy and energy demand has been clearly demonstrated in many studies (e.g. [1],[2]). This has led to the development of a range of models (principally [3], [4], [5], and [6], and recent published work such as [7] and [8]) for use in energy analysis that use the output from a standalone occupancy sub-model as the basis for determining the timing and magnitude of occupant-driven energy uses (e.g. lighting, appliances, hot water, etc.). Such models are used extensively in areas ranging from building stock modelling ([9]) to electrical systems analysis ([10]). These occupancy-driven demand models can be considered the current 'state-of-the-art'.

However, despite their utility and evident popularity, one criticism that can be levelled at the current generation of occupant-driven demand models is that they struggle to produce the distinct behaviours associated with individual occupants and households in real life. The reason for this is that the underlying basis for most such models is a simple low-order probabilistic model with a fixed mean and variance derived from a large dataset, which is called per timestep or event to identify actions. The output from this type of model will therefore converge to the average behaviour of the calibration population after a sufficient number of calls; although the speed of convergence is dependent on the model structure, with slower convergence possible with the addition of intra-day time dependencies, different day types etc ([11], [12], [13]). However, if the model results converge within the required analysis duration for each individual, this can severely limit the effectiveness of the output to mimic distinct occupant and household behaviours.

The rise of time-dependent renewables-linked microgeneration and the parallel reduction in the typical size of energy systems has made capturing individual behaviours more important in energy systems analysis ([8],[14]). The existing approach to modelling specific behaviours in areas such as building energy performance is to use models or analysis based on actual behaviours or demand data from small groups of households

*Corresponding author. E-mail address: graeme.flett@strath.ac.uk (G.Flett)

This is a peer-reviewed, accepted author manuscript of the following article: Flett, G., & Kelly, N. (2021). Modelling of individual domestic occupancy and energy demand behaviours using existing datasets and probabilistic modelling methods. Energy and Buildings, [111373]. https://doi.org/10.1016/j.enbuild.2021.111373

([15],[16]). However, this approach can often fail to recognise the 'Law of Small Numbers', where excessive relevance is placed on results generated from a very small and potentially unrepresentative sample [17]. While increasing the size of data collection studies to fully representative samples is typically impractical. A method is therefore required for a refined occupancy and demand modelling approach, which combines a probabilistic model that is representative of the overall population from which it is derived, but which has an improved ability to generate specific outputs that better match the diversity of behaviour seen in real households.

The rapid convergence in probabilistic models to the mean behaviour, and resulting limitations in reproducing individualistic behaviours, can be partially mitigated through an atomised probability model whose constituents are differentiated by factors such as age, employment status etc. However, a lack of calibration data limits the degree of differentiation. Therefore, the convergence risk within each differentiated sub-model remains. This paper assesses convergence in occupancy and appliance demand models used over extended durations (up to 1 year) for high resolution (sub-hour) occupancy and demand simulations. Improved modelling approaches are introduced that reduce convergence, through better representing diverse behaviours within each probability model, and crucially do not need more calibration data. The aim of the presented work is to identify the elements that exhibit rapid convergence, develop new methods and quantify the reduction in convergence to the mean compared to previous methods, and to assess the suitability for high resolution, stochastic occupancy and demand modelling.

Both the occupancy and demand elements of the principal, existing domestic-focused models are reviewed in Section 2. The previously developed occupancy and demand model by the Authors, on which the work presented in this paper is based, is also introduced. Section 3 expands on the aim of the presented work and explains the need for it in more detail. Section 4 explores the degree and speed of convergence in this existing model; to what extent it would impact on the ability to generate more realistic, diverse behaviours; and introduces the additional statistical manipulations applied to the existing occupancy and demand sub-models with the aim to reduce the convergence and generate more realistic behavioural diversities. Section 5 analyses the individual proposed manipulations on the occupancy and individual appliance sub-models and the overall impact on the creation of energy demand profiles that better represent individual household behaviours and demand patterns.

2. Existing Occupancy and Energy Demand Models

2.1. Existing Occupancy Models

The majority of existing high-resolution domestic occupancy models are based on two related, time-dependent, probabilistic modelling methods; Markov chains and discrete-event models. Markov chain occupancy models determine whether the occupancy state at time, t, is maintained or changes to another specific state at time, t+1, based on transition probability matrices (TPMs) calibrated using an underlying occupancy dataset. Markov chain-based domestic occupancy models were developed by Richardson et al [18] and Widen and Wackelgard [5], and have been further developed and utilised by a number of groups ([19], [20], [21]).

Discrete-event occupancy models determine both the time between occupancy transitions and the type of transition using probabilities. This type of model was developed by Wilke [6] and further enhanced by Aerts et al [22]. This dual determination per-event reduces the computational requirement, but was shown in [11] to have an inherent instability that is not seen in the more computationally intensive per-timestep Markov chain models.

A different method by Kleinebrahm et al [23] using neural networks has been developed that also attempts to solve the same individual behaviour replication limitations of Markov Chains that the presented existing and new work by the Authors' also addresses.

2.1.1. Previously Developed Domestic Occupancy Model

The approach to generating more statistically diverse occupancy and demand data presented in this paper uses a previously described occupancy model by the Authors [11], which uses a refined Markov chain approach with several enhancements to improve the ability to generate diverse but realistic household occupancy and demand data. These are: (1) significant occupant differentiation; (2) a 'higher-order' approach to occupancy duration prediction; and (3) accounting for the combined occupancy behaviours in multiple related occupant households (i.e. couples, parents and children). This enhanced approach was shown to perform better than the basic Markov chain models ([18],[5]) and the discrete-event approach developed by [6].

The Authors' model was calibrated using time-use survey data, which is a common data source for the majority of existing models ([5], [6], [18], etc.). Time-use surveys comprise activity diaries, including location, from multiple individuals and households with a sub-hour resolution. The primary limitation of time-use survey data is that the majority comprise single day diaries for each respondent. A significant number of individual diaries are therefore required to calibrate a probabilistic model with a sufficient depth of data, with the result that its underlying behavioural basis is a composite of the behaviours of a large number of people. The Author's model identifies three distinct occupancy states (sleep, active and out), which are then further differentiated to states such as TV-watching, working, and sleeping at or away from home, etc., as required.

The Authors' model uses multiple, separately calibrated, occupant and day-type transition probability matrix sets to reflect inherent differences in occupant characteristics between distinct populations and also between different types of days within populations. The tendency for convergence therefore needs to be analysed for both the individual matrix sets and the overall occupant models using the day-type sets in different realistic sequences. A more detailed description and review of the enhanced occupancy model is available in [11].

2.2. Existing Energy Demand Models

Occupancy-dependent, high resolution demand models require the integration of several separately modelled types of demand

(e.g. individual intermittently-used and continuously-powered appliances, standby power, hot water, lighting, heating, etc.). To capture the direct link between occupancy and specific energy demands, it is common practice to restrict the start time or overall duration of these demands to when a building is occupied ([24]). The effectiveness of the occupancy model in determining occupied periods therefore has a direct impact on the demand model performance. In addition to the occupancy driven element, the overall time dependency of specific demands is typically captured by time-dependent usage probability distributions, that are used to adjust the likelihood of use in different periods (e.g. to reflect increased probability of cooking appliance use at typical meal times). For example, Richardson et al [4] linked usage probability to related time-use survey activities and Widen and Wackelgard [5] applied appliance-specific distributions.

2.2.1. Previously Developed Domestic Energy Demand Model

The previously described demand model by the Authors, on which the presented new methods are based, incorporates both electricity and hot water sub-models. A full analysis and validation of the electricity sub-model is detailed in [12], with specific details and validation of the hot water sub-model presented in [13].

The electricity sub-model has several distinct modules to account for different types of demand, including lighting, TV-use, and three separate modules for intermittently-used appliances to account for different use characteristics. The hot water sub-model uses the same basis as the intermittent appliance use module applied to high-use, short-term need driven appliances (e.g. kettles, microwaves).

The developed model differs from existing models most significantly in relation to the intermittently-used appliance modules. Existing models (e.g. [4]) typically use a first-order per-timestep probability approach, with each timestep analysed sequentially to determine if an appliance use event occurs. This was shown in [12] to generate unrealistic sequencing as the first-order approach does not consider time since a previous usage event. The developed model substitutes a higher-order discrete-event approach, where the number of uses on a specific day is first determined and then the timing for each individual use event predicted based on occupancy and time-specific use probabilities. This discrete-event approach was shown to be more accurate in determining the overall distribution of uses for individual households, as it allows use timing to be individually calibrated based on the xth use in a n-use day.

The convergence analysis and demand model enhancements presented in this paper are focused on these intermittently used electrical appliance and hot water modules. All are strongly occupancy-driven, have significant time dependency, and where use frequency and time-of-use can vary significantly between individual households.

3. Aim

The aims of the presented work were to firstly, identify critical areas of rapid and unrealistic convergence to average calibration population behaviour in previously developed probabilistic occupancy and demand models that would limit their effectiveness for types of energy system modelling where individual behaviours have a significant impact. Secondly, to identify statistical manipulations that could be applied to reduce convergence or to further factor the output to individual behaviours, without requiring significant changes to the model structure or new calibration data, and which does not reduce the statistical consistency of the overall model to the calibration data. And finally, to quantify the reduction in convergence and improvement in the simulation of individual household behaviours, and the suitability for stochastic modelling of occupancy and demand, for the new individualised model basis.

4. Model Convergence Analysis and New Method Development

As outlined in Section 1, output from first-order stochastic probability models will eventually converge to the underlying mean behaviour after a sufficient number of calls without further statistical manipulation. In this section the occupancy and demand models previously developed by the Authors ([11],[12],[13]) are analysed to determine to what extent the speed of convergence prevents individual behaviours being captured and describes the further statistical manipulations developed to address this.

4.1. Occupancy Model

Convergence behaviour is considered at two levels: (1) within individual occupant and day-type Markov-chain transition matrix sets (each set comprising multiple matrices differentiated by time of day for the same occupant and day type (weekday/weekend, working/non-working)); and (2) for the overall occupancy model output when multiple day-type matrix sets are used in sequence over longer periods for individual occupants.

4.1.1. Speed of Convergence of Individual Occupancy Models

Each Markov-chain transition matrix set is generated from the single-day activity diaries of a large number of individuals. For each set, the average behaviour will therefore converge to that of the calibration population; an inherent feature of all Markov-chain models.

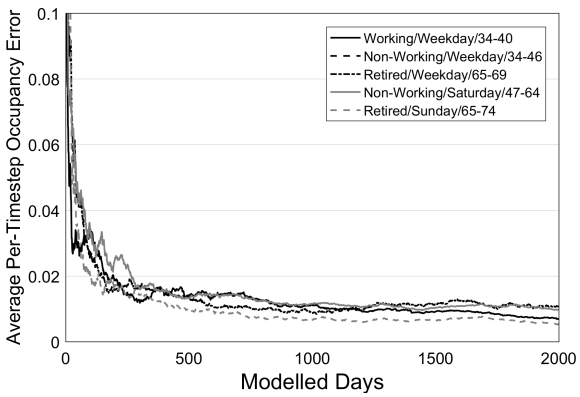

Fig. 1. Occupancy sub-model convergence - average per-timestep 'error' to calibration dataset average by number of modelled days.

Convergence in this case is determined by the average per-

timestep difference ('error') between the calibration population active occupancy and the model output. Analysis has shown that each set converges rapidly to the average behaviour over approximately 100-200 modelled days as shown in Figure 1 for selected examples.

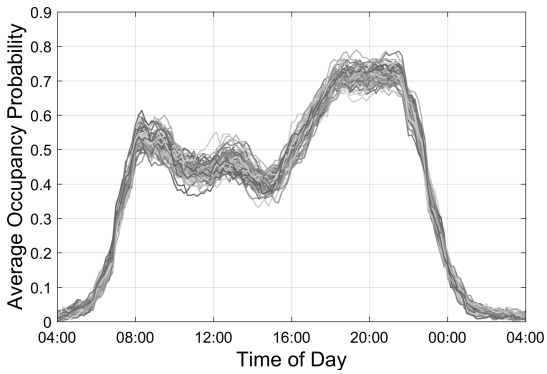

(a) 'Richardson' model

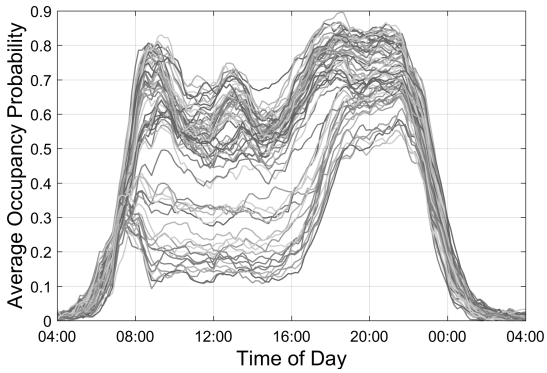

(b) 'Composite' model

Fig. 2. Occupancy sub-model average per-household active occupancy results for 70 single-person households for annual duration model runs.

4.1.2. Impact of Convergence for Overall Model Performance

Occupancy models typically combine different occupant and day type modules in specific sequences to simulate individual lifestyles. Day type differentiation ranges from simple distinctions between weekdays and weekends to more detailed differentiation, including by age range and employment status. Further analysis is therefore required to determine if the use of multiple occupant and day types in different realistic combinations, allows individual occupancy behaviours to be captured despite the convergence of each day type sets.

The Richardson et al model [18] differentiates by total occupant number and day type (weekday/weekend). The Authors' previously developed occupancy model [11] has significant differentiation by occupant type, age, and employment status on the modelled day, and realistic individually tailored day-type sequencing reflecting actual variation in typical working day patterns. Both models have been used for analysis with the first-order (i.e. only the current occupancy state is considered) model based on [18] referred to as the 'Richardson' model and the Authors' higher-order (i.e. the duration of the current state is considered) model referred to as the 'composite' model. The 'Richardson' model output used for analysis is based on a single-person household with weekday/weekend differentiation only. The 'composite' model output is based on a nationally representative set of single-person household characteristics. This comparison allows a judgement on the impact of increased occupant differentiation.

There is no UK-specific large-scale occupancy dataset that tracks individual occupancy for longer than one day that would allow the overall long-duration model convergence to be assessed. However, this assessment can be made by comparing the results with the Dutch 2005 TBO Time-Use Survey (Dutch TUS) [25], which includes seven-day diaries. It is probable that one-week diaries for an individual will exhibit different behaviour variation than if the same individual was monitored for a longer period, and Dutch occupancy behaviours may vary from equivalent UK behaviours, but the dataset allows for an initial judgement to be made on model replication of specific occupant behaviours.

Two specific types of convergence need to be considered; (1) the convergence in the timing of key occupancy transition periods (waking, going to sleep, etc.) and (2) convergence in the average occupancy per-timestep. Figure 2 shows the time-dependent average annual occupancy profiles for 70 single-person households for both the 'Richardson' and 'composite' models. The 'Richardson' model output, with no further occupant differentiation, shows significant uniformity for both average occupancy (vertical variation) and key occupancy transition timing (horizontal variation). For the 'composite' model, with additional differentiation for age and employment status and a representative set of occupants, there is greater variation in average occupancy profiles, however, the key transition timing remains unrealistically convergent.

To demonstrate that the 'composite' model key transition timings are unrealistic, Figure 3 shows the waking time mean and standard deviation (in minutes) distribution for retired single-person households from the 15-minute resolution Dutch TUS data and the equivalent for the 10-minute resolution 'composite' model based on 100 one-week and 100 one-year duration runs. The TUS data shows a range of behaviours, including zero standard deviation results. (A zero result implies that the person transitioned within the same 15-minute period on all monitored days.) The equivalent occupancy model results show a more random distribution centred on the average behaviour for the one-week models and a highly converged distribution for the one-year models; both with a higher average standard deviation than the TUS data. This highlights that the 'composite' model output is overly variable per modelled event for each individual but overly convergent, in general, in comparison with real behaviours.

To assess whether the average occupancy variation shown in Figure 2 (b) for the 'composite' model is realistic, further analysis of overall average occupancy within each differentiated subgroup was undertaken, as shown in Figure 4, for one-week and one-year duration models in comparison with the Dutch TUS data. This indicates that for one-week duration models, the output is slightly less variable that the TUS data. However, for

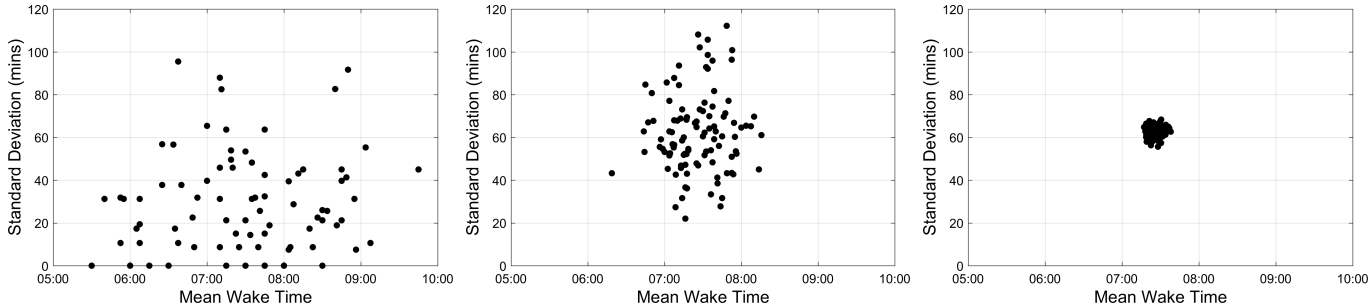

(a) Dutch TUS **(b)** 'Composite' model - 1 Week **(c)** 'Composite' model' - 1 Year

Fig. 3. Mean and standard deviation of wake time per occupant for single-person, retired households for the Dutch 2005 TBO TUS dataset [25] and 'composite' model output.

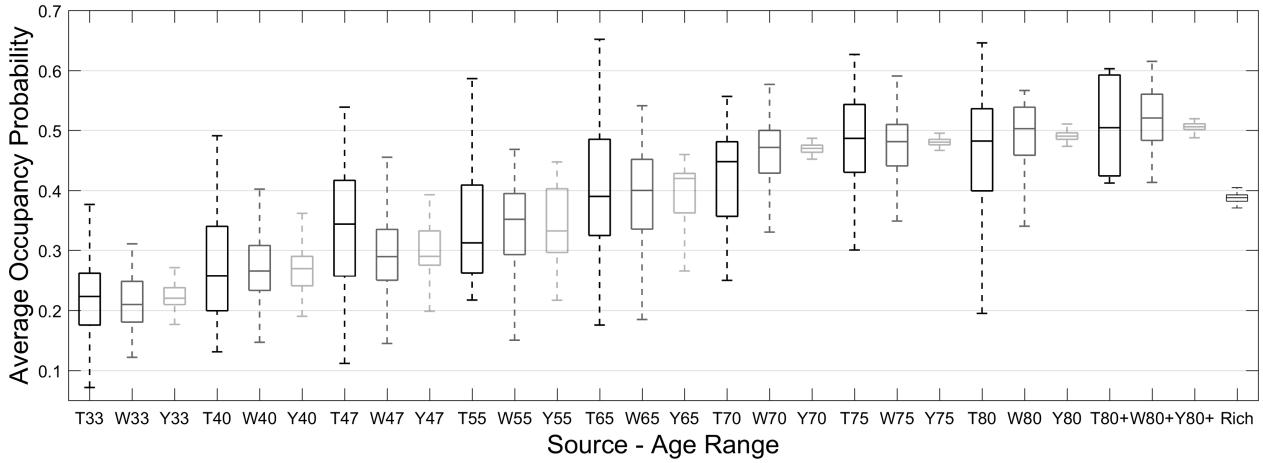

Fig. 4. Per-individual active occupancy range comparison for the 1-week Dutch 2005 TBO TUS dataset [25] and equivalent higher-order Markov chain occupancy model output. (T=TUS, W=1-Week Model, Y=1-Year Model, Number=Upper Age, 'Rich'=Richardson et al) (Box=25%/50%/75%, Limits=0.3% and 99.7%)

one-year models; the reduced variance is more marked, particularly for the retired householders which lack the different combinations of working and non-working days for other age groups. This indicates that the 'composite' model does not capture the normal variations in average active occupancy within each population.

4.1.3. Reducing Occupancy Transition Timing Convergence

To reduce the observed transition timing shortcomings, the Authors' occupancy model has been modified to better capture individual transition timing characteristics while maintaining consistency to the overall population behaviour. To achieve this, the Dutch 2005 TBO TUS (Dutch TUS) dataset [25], which includes seven-day diaries, has been used to calibrate the model for individual behaviours relative to the average behaviour for each defined calibration group. It is assumed that the variability in this dataset, if not the specific timings, are representative of any developed country population. It is acknowledged, as outlined above, that one-week diaries may not be sufficient to accurately capture long-term behaviours, and that equivalent UK-specific data over a longer period would significantly enhance the proposed method. However, this is not available at the time of writing.

Analysis of the Dutch TUS data has shown that transition time behaviour within each household-type varies significantly between respondents. This can be demonstrated by comparing the mean and standard deviation (in minutes) of the relevant transition times. As demonstrated in Figure 5 for the waking times of different single-person householder groups, there is significant variation in the mean timing, with some respondents showing consistent patterns of behaviour (low standard deviation), others are more erratic. This behaviour variation is replicated across all populations for both wake and sleep transition timing.

To capture this behavioural variation, Markov chain occupancy probability matrices are assumed to provide two levels of behavioural information. The first level is the direct time-specific detail used by existing models, including the previously developed 'composite' model, to probabilistically generate statistically consistent, if unrealistically convergent, stochastic outputs. The proposed new method assumes there is also an underlying level of occupancy behaviour information inherent in Markov chain occupancy calibration data that captures typical transition patterns that are relative to specific actions and therefore period- rather than time-specific. For example, there is a relationship between the time a person wakes and leaves the house that is, to a degree, independent of the specific timing and related to the typical duration between the two events.

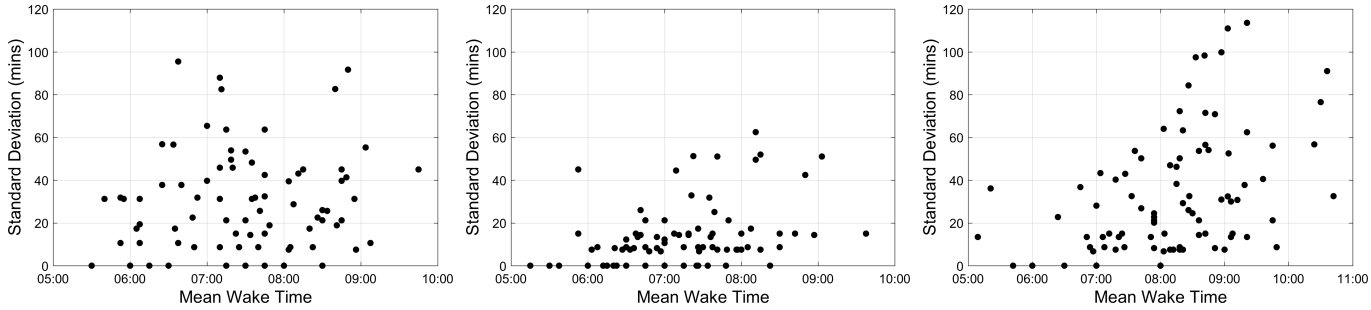

(a) Retired **(b)** Working **(c)** Non-working

Fig. 5. Wake time mean and standard deviation for retired, working and working age non-working single-person householders from the Dutch 2005 TBO TUS dataset [25].

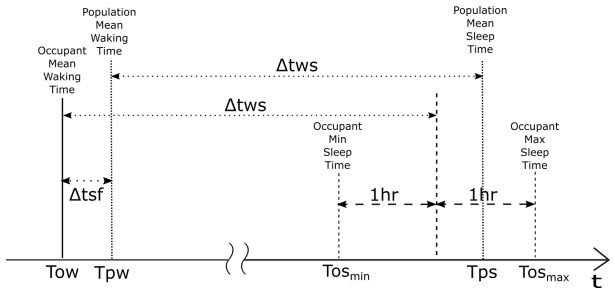

Fig. 6. Graphical representation of process to convert TUS average wake and sleep times to individual-specific equivalents.

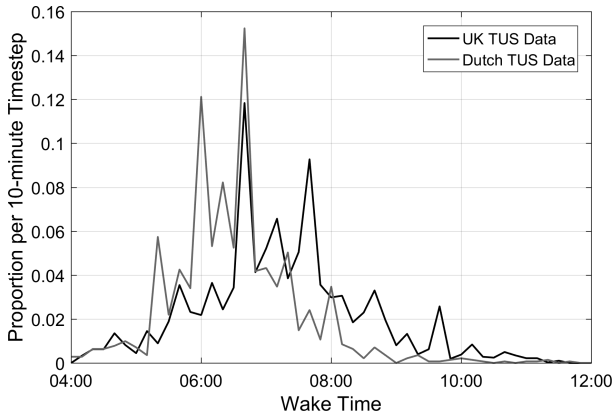

Fig. 7. Proportion of weekday wake times per 10-minute timestep from the UK 2000 [26] and Dutch 2005 TBO TUS [25] dataset single-person householders.

The 'composite' occupancy model uses a higher-order basis that accounts for the time since the last occupancy transition (based on the three states: sleep, active and absent). Therefore, this period-action rather than time-specific behaviour is also partially captured within the existing calibration matrices and model structure, although the convergence analysis has shown that this is insufficient to capture individual behaviours. The 'composite' model approach has therefore been further manipulated to capture the underlying inherent period-action behaviour without impacting the overall statistical basis or introducing unrealistic individual occupancy patterns. A modified version of the 'composite' model, hereafter referred to as the 'individualised' model, has been developed to manipulate the transition timing basis for each individual modelled. This has been achieved by modifying the time structure of the Markov model for individuals to account for individual behaviours relative to the population average.

4.1.4. New Occupancy Transition Timing Method

The following section outlines each additional manipulation to the previously developed 'composite' occupancy model [11] included in the newly developed 'individualised' occupancy model. Figure 6 shows the overall process and each of the manipulations required graphically.

Occupant Mean Waking Time (Tow) – Analysis of the Dutch TUS data shows that there is no significant difference in the average variance per person when waking and sleep times are directly compared. Waking time has therefore been arbitrarily selected as the anchoring statistic for the revised model. Each individual has been allocated an average waking time based on the probability distribution for the equivalent UK TUS population. (Figure 7 shows this distribution for single-person householders). It is assumed that some of the early and late times observed in the UK TUS data are outliers for those individuals and a comparison between the one-week Dutch data and single-day UK data suggests a 5-10% greater range in the single-day UK data, therefore the potential average wake times are restricted to those in the middle 90% of the UK TUS distribution.

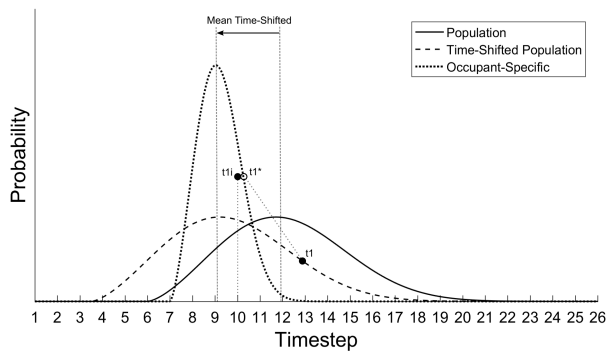

Fig. 8. Graphical representation of a calibration population distribution conversion to individual time-shifted basis with lower variance.

Waking period time-shifting – The 'waking period' is defined as 3am to 10am. For each calibration population, the population mean waking time is determined. For each individual, the selected time-dependent Markov chain probability data in this 'waking' period is time-shifted (Δtsf) based on the difference between the individual's mean waking time and the relevant population average (*Tow* and *Tpw* respectively in Figure 6).

Sleep period time-shifting – Analysis of the Dutch TUS data shows an average waking duration of 17.25hrs (Δtws), with the majority of individuals being uniformly distributed by +/- 1hr of this value. This range closely matches the average waking duration for the unmodified 'composite' model. The sleep-transition period Markov chain probability data (9pm to 3am) is therefore time-shifted for the individual, in the same manner as the 'waking' period, by a randomly selected amount that is +/- 1hr of the wake period time-shift (i.e. between Tos_{min} and Tos_{max} in Figure 6).

Variance factor – The output from the time-shifted Markov-chain process without further modification maintains the overall timing variance of the composite population behaviour, which has been shown to significantly exceed individual variance for almost all individuals (see Figure 3). For example, the single-person working day population has an average waking time standard deviation of 60.3 minutes and the equivalent for the retired single-person population is 73.6 minutes. The equivalent for each individual in the Dutch TUS dataset are 28.0 and 20.6 minutes respectively, with only 8% and 4% of each population exceeding the population variance. To achieve realistic individual behaviour variance, each modelled individual is randomly allocated a standard deviation for each defined transition based on the equivalent Dutch TUS dataset occupant-type distribution. The difference between the individual-specific standard deviation and the calibration group average is then used to further alter the Markov chain timing basis.

Figure 8 demonstrates the rebasing method graphically, with the individual-specific transition timing probability distribution with the rebased mean (left-shifted = earlier than average) and standard deviation (narrower = lower variance than average) shown in comparison to the population average (with both the original and time-shifted versions shown).

The required rebasing manipulation is achieved by first running the time-shifted population Markov chain algorithm without any further manipulation until a relevant transition is identified (if no transition is identified the model proceeds as normal to allow for the probability of uncharacteristic behaviours). The cumulative distribution function (cdf) value for the transition point (*t1*) is determined and the equivalent *cdf* value (*t1**) on the individual-specific variance distribution identified. The model timestep is then reset to the closest integer timestep (*t1i*) to the identified point *t1** and the relevant transition set at this timestep. If *t1i* is before *t1*, the model deletes all modelled timesteps after *t1i* and resets the model timestep to *t1i*. If *t1i* is after, the unmodelled timesteps up to *t1i* are set to the preceding state and the model continues from *t1i*.

Other Key Transition Periods – The same process is also used for the key morning leaving and evening returning transitions, if such a transition is predicted within defined periods (morning leaving - from waking until 2.5 hrs after mean waking time; evening return - ± 1.5hrs of mean return time). If these transitions are not predicted, the model proceeds without further manipulation.

4.1.5. Residual Average Occupancy Convergence

The inclusion of significant occupant and day type differentiation (see 4.1), and also realistic occupant work weeks within the occupancy model, ensures a degree of variation in average occupancy, both within and between individual type populations. However, Figure 4 indicates that, while this differentiation results in improved performance over existing methods, the occupancy model output remains less variable than real behaviours.

Addressing average occupancy convergence is less straightforward than for transition timing convergence, as it is an underlying function of the overall probability model rather than related to specific transitions in distinct time periods. However, the main determinant of average occupancy is the balance of 'active' and 'out' periods, with 'sleep' duration variability addressed by the transition timing method identified in the preceding section. Relative average occupancy can therefore potentially be manipulated to account for variable in-population behaviour by manipulation of the 'active-out' transition probabilities. A method to address this is in development but is not yet integrated within the overall model and therefore not covered in this paper.

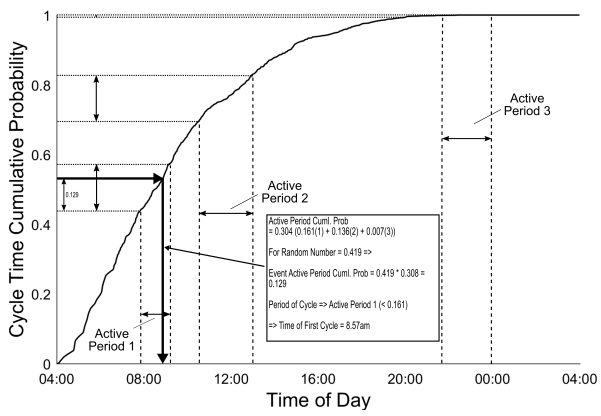

Fig. 9. Use (cycle) time prediction model example.

4.2. Domestic Electrical and Hot Water Demand

4.2.1. Individual Appliance Model Convergence

The Authors' previously developed demand model [12] combines the output from the occupancy model with an event-based probability approach to identify individual appliance and hot water use times that are differentiated by household type, daily use number, and daily use consumption. Appliance use timings from the 2010/11 Household Electricity Survey (HES) dataset [27] and hot water use timings from the 2008 Energy Savings Trust dataset [28] where analysed and converted to cumulative probability density distributions, which were further adjusted to account for relative levels of occupancy per time period. For

each modelled day the use timing is further restricted to occupied periods only as addressed in detail in [12]. Figure 9 shows an example of a cumulative use timing probability distribution and the occupancy period restrictions (see 'Active Periods'), with the occupancy in this new individualised demand method being generated from the 'individualised' occupancy model basis. These distributions, however, do not account for any household-specific or habitual behaviours related to the time-dependent use of electrical appliances and hot water.

Analysis of the predicted use start times for each modelled household can be made by converting the identified start times to the relevant cumulative density function (cdf) value between 0 and 1. The mean and standard deviation of these cdf values per household gives a measure of the specific household usage behaviour compared to the average behaviour. For example, a household mean start time cdf value significantly lower than 0.5 indicates use that is typically earlier than the population average. The lower the standard deviation, the more consistent the timing of each use (i.e. a low std. dev. value indicates a use that is consistently at the same time each day and high value a more random use behaviour).

Analysis of the HES and EST datasets has shown that significant use timing behaviour differences exist for most intermittent demands. However, in terms of both likelihood and overall impact on demand, is most significant for the following; kettles, washing machines, dishwashers, cookers/ovens, and the higher volume (>15 litres) hot water events most associated with baths and showers. Assuming average behaviour should therefore, to some degree, impact model accuracy (this was confirmed by comparative analysis detailed in 5.2 and 5.3).

Figure 10 (a), (d), and (g) shows the range of household *cdf* mean and standard deviation values for washing machine and cooker cycles for all HES dataset households and 'Very High' (30 litres+) hot water cycles for all EST dataset households. These results indicate that there are significant behavioural differences between households and that this is not merely a function of the number of observed events (indicated by the size of each data point). Figure 10 (b), (e) and (h) show the equivalent distributions for households modelled using the 'composite' model, indicating a strong tendency for convergence to the average behaviour. The results also demonstrate that the convergence tends to increase with the number of modelled events, as would be expected for a highly probabilistic model. The 'composite' model behaviour is consistent for all the specific demands listed above.

4.2.2. New Individual Appliance Timing Method

Appliance and hot water use timing is simulated by converting a generated random number to a time based on the relevant start time cumulative probability distribution, with the potential start times limited to occupied periods, as shown in Figure 9. The new 'individualised' model basis restricts the random number generation to replicate the realistic distribution of appliance use timing behaviour per household.

The 'individualised' random number generation manipulation is achieved by reducing the ascending range of use timing *cdf* values for each HES or EST household into 21 representative quantiles and converting each quantile value into one of ten *cdf* value ranges (1=0-0.1, 2=0.1-0.2, etc.). The *cdf* range transitions between each of the quantiles were determined for each household and used to calibrate a 21-element sequential Markov chain model for each specific occupant-driven demand.

The Markov chain model generates new *cdf*-range distributions for each modelled household based on a probabilistically assigned midpoint (11th) quantile, with the Markov chain model working in both directions from the midpoint to the minimum (1st) and maximum (21st) quantile values to allow the critical midpoint value to be further factored based on household occupancy timing distribution compared to the average. Four typical resultant distributions for households that exhibit distinct use behaviours (early/average/late with low variance, and average with high variance) for the cooker module are shown in Figure 11. For example, for early use behaviour with low variance ('Early/Low Var'), the random numbers are restricted to the 0 to 0.2 range (ranges '1' and '2'), while average use behaviour with high variance ('Average/High Var') shows a relative linear distribution across all ranges.

Within the 'individualised' cycle start time module, a value is selected randomly from the household-specific 21-quantile distributions and then the actual value used is selected randomly within the range (e.g. a 'CDF Range' value of 4 randomly selected from the distribution is converted to a number randomly selected between 0.3 and 0.4). The determined value replaces the original 'composite' model equal probability random number generation used to determine the cycle start time, in order to skew the generated start times to reflect individual household behaviours.

The use of a relatively small number of quantiles and ranges ensures that the broad overall pattern of potential behaviours is captured but that there is sufficient variation from randomly selecting values within the ranges to ensure that the model is not forcing close replication of the input data. This 'quantile' method was preferred to other methods of skewing probability distributions as it allows for multiple and well separated periods of higher use probability.

A proportion of the observed appliance and hot water use timing variation is assumed to be the result of occupancy variations. In the absence of data directly linking occupancy with occupant-driven energy demands, the mean timing value for each household is modified by the extent to which the average household occupancy is earlier or later than the average population behaviour. The 'individualised' method as currently implemented does not, however, account for specific daily occupancy patterns. For the original 'composite' method, the potential cycle times are first limited to the occupied periods and then the specific time is determined based on the generated value between 0 and 1 which is used to locate it proportionally within the occupied period. A similar process is used for the 'individualised' method, with only the random number restricted to a smaller range, therefore the behaviour is only skewed within the general timing behaviour characteristics of the population-level cdf distribution and not forced to specific times.

Further improvement of this method to account for highly distinct use patterns by linking use to specific time periods is re-

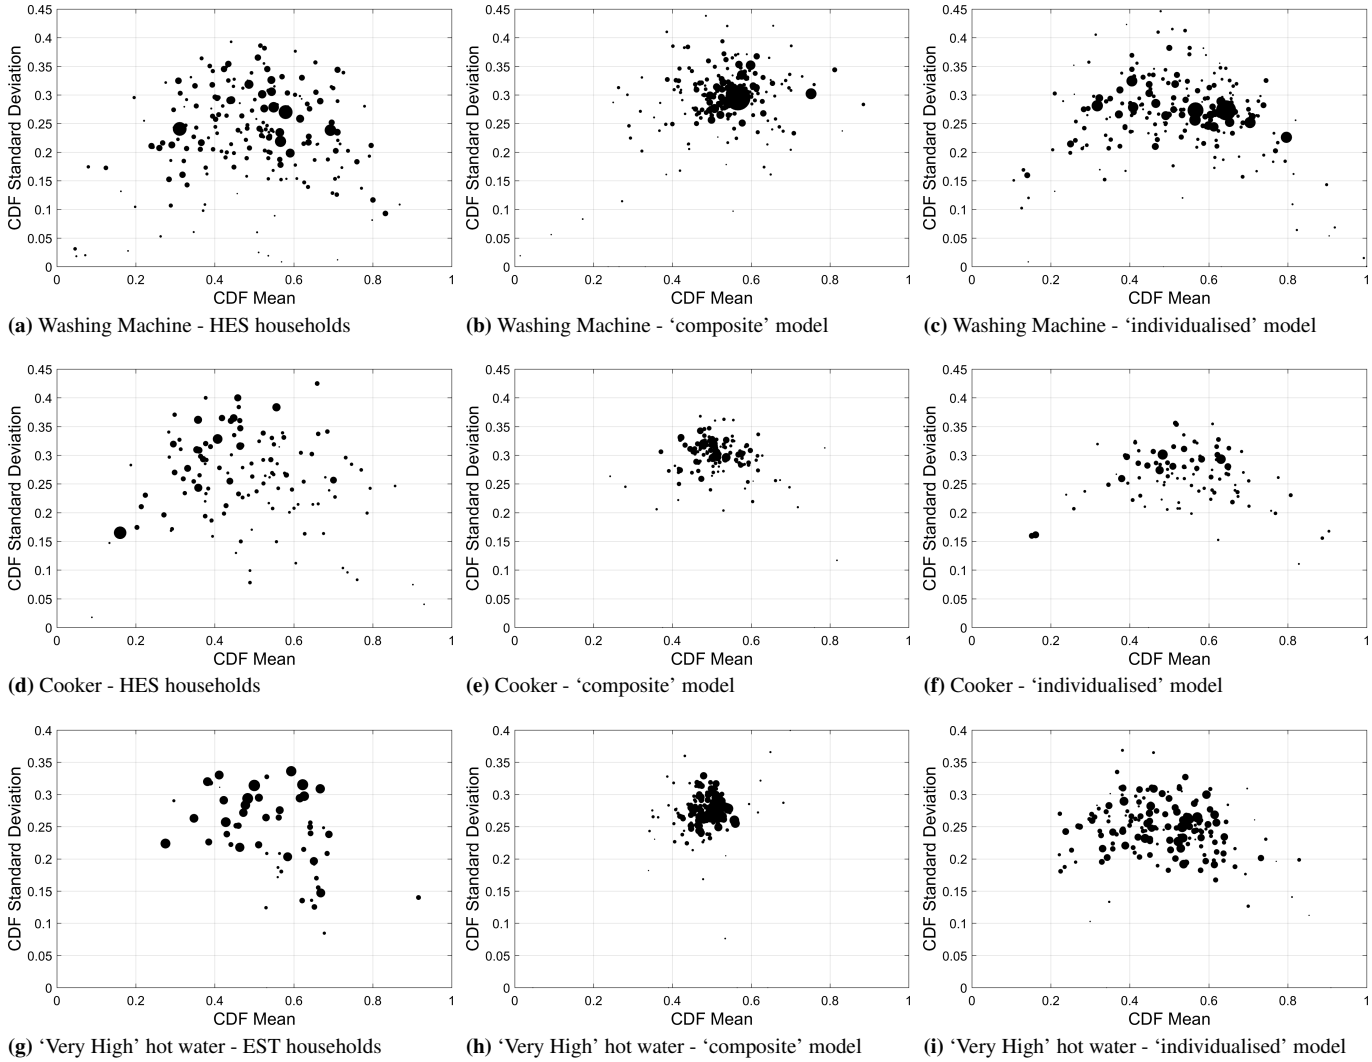

(a) Washing Machine - HES households **(b)** Washing Machine - 'composite' model **(c)** Washing Machine - 'individualised' model

(d) Cooker - HES households **(e)** Cooker - 'composite' model **(f)** Cooker - 'individualised' model

(g) 'Very High' hot water - EST households **(h)** 'Very High' hot water - 'composite' model **(i)** 'Very High' hot water - 'individualised' model

Fig. 10. Cycle start time cumulative probability function mean and standard deviation per household. Data for the 'HES households' distributions from [27] and for the 'EST households' distribution from [28].

quired for better replication of applicable households. In particular, an integrated occupancy and demand dataset (as opposed to the distinct datasets underpinning the work) would allow the statistical relationship between occupancy and cycle timing for individual households to be better incorporated.

5. Individualised Method Evaluation

5.1. Occupancy Transition Timing Model Evaluation

The following section analyses the impact of the new occupancy modelling method outlined in 4.1.4 that aims to provide closer replication of individual occupancy behaviours.

5.1.1. Convergence Analysis

Average occupancy results using the 'individualised' model basis are shown in Figure 12 for the same population as the previously developed 'composite' model results shown in Figure 2. This demonstrates that the new 'individualised' approach generates significantly more variation between individual single-person householders in the timing of the key transition periods (waking, morning leave, afternoon return, sleep), while maintaining the same general occupancy patterns.

Analysis of the 1-year duration results for a number of the distinct occupant-type models was undertaken to determine the detailed performance of the 'individualised' approach. For the 74-79 year old single householder group, it was determined that approximately 20 annual runs of the 'individualised' model was required on average to achieve the same level of convergence as a single equivalent run of the original 'composite' model. Similar values were found for all other populations. The results show that the individualised model retains the overall convergence to the population average but with significantly more per-individual diversity.

5.1.2. TUS Dataset Replication

A direct comparison between the model output and individual occupant behaviour over an extended period is restricted by the lack of long-term occupancy data. However, a statistical com-

Table 1
Occupancy model validation metric comparison for three single-person householder populations for the 'composite' and 'individualised' models based on 1000 annual model runs.

	Per-Run Average				Overall - 1000 Runs			
	AO_Conv (x E-3)	DurDist Sleep	DurDist Active	DurDist Out	AO_Conv (x E-3)	DurDist Sleep	DurDist Active	DurDist Out
Working 34-40 - 'Composite'	23.0	1.39	0.66	2.37	20.1	1.15	0.44	1.70
Working 34-40 - 'Individualised'	87.4	4.56	1.59	4.95	26.8	0.87	0.48	4.13
Non-working 34-46 - 'Composite'	27.6	1.67	0.81	1.25	19.3	1.18	0.38	0.81
Non-working 34-46 - 'Individualised'	40.2	5.08	1.46	2.05	18.0	1.95	0.75	1.39
Retired 70-79 - 'Composite'	15.0	1.08	0.78	0.94	10.2	0.44	0.28	0.60
Retired 70-79 - 'Individualised'	119.9	6.59	2.11	1.16	28.1	4.47	1.32	0.54

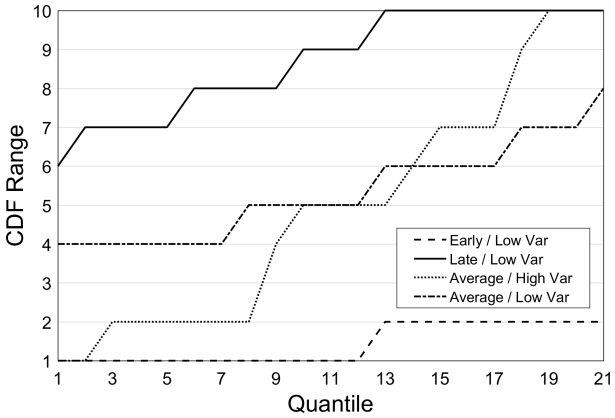

Fig. 11. Example individual household cycle start time cumulative distribution function (cdf) quantile distributions from developed Markov chain model.

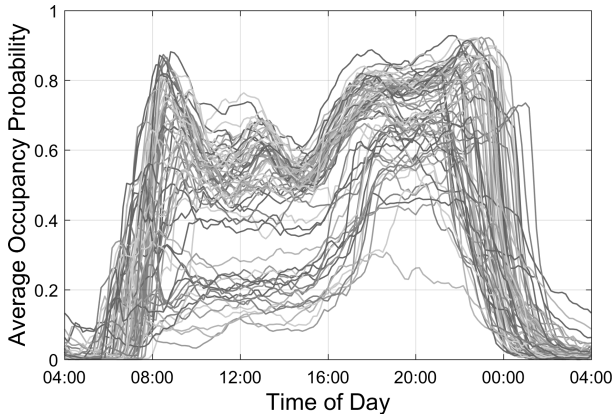

Fig. 12. Occupancy sub-model average per-household active occupancy results for 70 single-person households for annual duration model runs for the 'individualised' model basis.

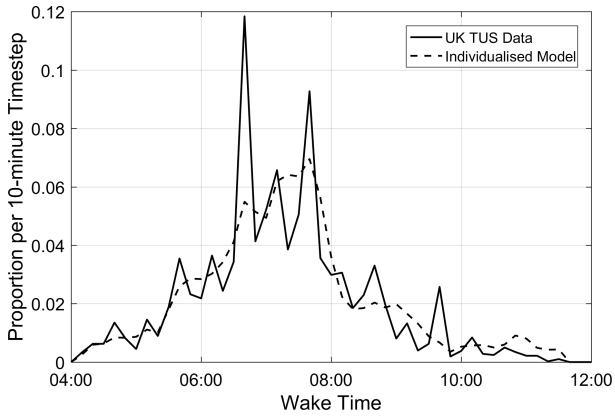

Fig. 13. Wake time comparison between UK 2000 TUS dataset [27] and 'individualised' occupancy model output for single-person householders.

parison between the TUS data and the occupancy model output for each occupant-type group is possible. The 'individualised' occupancy approach should replicate the timing behaviours of each equivalent TUS population for the identified transitions. Replication analysis for the waking period (see Figure 13) for all single-person households shows good correlation between the UK TUS distribution and model results. A smoother distribution is to be expected for the model results as they include significantly more data points than the TUS dataset (25550 vs 1159). Similar correlations are observed for other transitions and occupant-type groups.

The 'individualised' model has been calibrated to reflect the average waking times from the UK TUS dataset, therefore exact replication of the Dutch TUS distributions shown in Figure 5 is not expected. However, the replication results for 100 annual-duration retired household models using the 'individualised' model shown in Figure 14 highlights the improvement in replicating the variance in behaviour compared to the tight convergence shown in Figure 3(c). The model distribution matches the characteristics of the Dutch TUS data equivalent, with a slightly later average reflecting the overall behaviour difference shown in Figure 7.

Further validation of the 'individualised' method was undertaken using the metrics defined in [11], where full details of their derivation can be found. These are described briefly below and it is sufficient to note that the smaller the metric result the better the replication in all cases. This validation was performed for three single-person weekday household model groups (working 34-40, non-working 34-46 age and 70-79 age range). The results are shown in Table 1.

Average Occupancy Convergence Metric – determines the average per-time step occupancy error between the Time Use Survey (TUS) input data and average model output for each

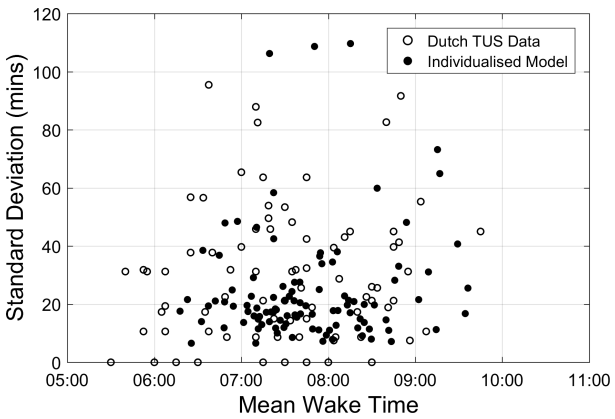

Fig. 14. Wake time statistical comparison between the Dutch 2005 TBO TUS dataset [25] and 'individualised' occupancy model output for single-person, retired householders.

occupancy state - quantifying the quality of calibration of the model. Equation (1) below is based on 144 data points per day (10 minute time steps).

$$AO_Conv_{state} = \sum_{t=1}^{144} \frac{\left|\bar{P}^{mod}_{state}(t) - \bar{P}^{tus}_{state}(t)\right|}{144} \tag{1}$$

where, AO_Conv$_{state}$ is the Average Occupancy Metric for state, *state*, $\bar{P}^{mod}_{state}$(t) is the average modelled probability for state, *state*, at timestep, *t*, and $\bar{P}^{tus}_{state}$(t) is the average probability for state, *state*, at timestep, *t*, derived from the input Time-Use Survey data.

For the *AO_Conv* metric, the per-run average for the 'individualised' model shows significantly more variation, as expected for a method that replicates individual behaviour variance, but the overall convergence after 1000 annual runs is similar for each of the populations, confirming that the 'individualised' method does not significantly reduce overall statistical similarity to the calibration dataset. However, when the error per timestep is analysed an apparent period of weaker replication is seen in the sleep transition period, particularly for the retired population. This suggests that the simple correlation between wake and sleep time used requires a more complex statistical basis for improved accuracy.

State Duration Distribution Metric – (hereafter referred to as DurDist) is used to assess the ability of a model to generate a realistic range of occupancy state durations. It compares the difference in the cumulative probability function (CDF) at each 10-minute duration range for the histograms of the model generated results and TUS data in order to determine if the generated occupancy profile replicates the occupancy state durations seen in the TUS. The 'error' is the sum of the absolute difference between the model and TUS data CDFs at each duration value for each state (see Equation (2)).

$$DurDist_{state} = \sum_{d=1}^{144} \left| \sum_{d=1}^{d} \bar{P}^{mod}_{state}(d) - \sum_{d=1}^{d} \bar{P}^{tus}_{state}(d) \right| \tag{2}$$

where, $\bar{P}^{mod}_{state}(d)$ is the probability of a modelled state duration of *d* for state, *state* and $\bar{P}^{tus}_{state}(d)$ is the probability of a state duration of *d* for state, *state*, derived for the input Time-Use Survey data.

For the *DurDist* metric, the performance is broadly similar, with higher variation seen for individual model runs but broadly similar performance overall.

There were two areas of weaker performance for this measure, sleep duration for the retired population and the absence duration for the working population. The absolute impact of the error in each case was small (approximately 45 minutes) but suggests that more complex statistical relationships are required to link wake and sleep times and to link leave and return times for certain populations to improve modelling.

The results show that the modifications to the model do not significantly impact on the statistical consistency, although the performance is weaker in some areas as a result of the manipulations to the underlying Markov chain characteristics. Analysis of the overall impact on energy demand modelling accuracy (see Section 5.3) is required to make a final assessment on whether the benefits of the increased variation and improved individual behaviour replication outweigh any reduction in statistical consistency.

5.2. Individual Appliance Use Timing Model Evaluation

The following section analyses the impact of the newly developed individual appliance use timing method outlined in 4.2.2 that aims to provide closer replication of individual household use behaviours.

Visual and statistical analysis of the *cdf* mean and standard deviation values from the 'individualised' model results in Figure 10 (c), (f), and (i) shows an overall distribution that is significantly closer to the measured data ((a), (d), and (g)) than the 'composite' model results ((b), (e), and (h)). The results also indicate that there is no evidence of greater convergence to the mean behaviour for households with a higher number of modelled use events (each data point is scaled), which was a critical identified performance problem for the 'composite' model. The residual discrepancy, particularly in capturing extreme behaviours, results from the lack of direct integration between the occupancy and use event timing sub-models, as a result of current data limitations, that does not easily allow habitual (i.e. low standard deviation) behaviour to be closely replicated.

A measure, the Earth Mover's Distance (EMD), is used to determine the similarity in the cdf mean and standard deviation distribution, including inherent factoring for the number of use events included in the data for each point in the distribution to specifically account for the tendency for convergence to the mean behaviour with an increasing number of modelled uses in the 'composite' model basis. The EMD method effectively considers one distribution as equivalent to a series of mounds of 'earth' and corresponding 'holes'. It characterises a particular profile, by quantifying the minimum 'work' required to fill the holes in terms of moved 'earth' and distance to the 'holes'. The stochastic nature of the overall demand model means that exact replication, as judged by EMD, of the HES and EST distributions is not expected or desirable. However, it does pro-

vide a more robust assessment of the degree to which the 'individualised' model improves specific behaviour replication and prevents increasing convergence with the number of modelled uses.

Table 2
Average 'Earth Mover's Distance' measure for different specific demands. Dataset data for analysis from [27] and [28].

	Cooker	Washing Machine	Hot Water 'High'	Hot Water 'Very High'
'Composite' Model	7.31	6.42	3.23	5.19
'Individualised' Model	4.93	3.14	2.06	3.22

There is a significant increase in distribution similarity (lower EMD) for the 'individualised' model as shown in Table 2. The results also show that for the composite model the results per run are less variable. For example, for the cooker they only range from 6.82 to 7.70, while for the 'individualised' model they range from 3.51 to 7.86, indicating that the 'individualised' method also achieves the broad aim for achieving increasingly realistic per-run stochastic variability while retaining overall underlying convergence after a significant number of runs.

Whilst the individual cycle start time method can be further improved, it has been shown to be a significant improvement on the 'composite' model basis for each specific demand. Further consolidated analysis is required to determine if the method improves the overall demand model performance and to determine how applicable it is for energy system development. This analysis is presented in the following section.

5.3. Impact on Overall Occupancy and Demand Model Performance

5.3.1. Electricity Demand

The piecewise aggregate approximation with Euclidean distance (PAA-ED) similarity method introduced in [11] can be used to compare the overall electricity demand sub-model output (combining individual appliance and lighting outputs), with and without the defined individualised occupancy and cycle behaviour modifications, to actual measured household data. The PAA-ED method simplifies each 144-timestep average demand profile to a 36-time segment approximation based on ranges of demand, which can be compared using a standard Euclidean edit distance similarity measure, with a lower value indicative of greater similarity. Two distinct PAA-ED measures are used: 'Timing' which reflects the variation in relative timing of demand but ignores differences in absolute demand; and 'Overall' which captures both differences in relative timing and absolute demand.

Based on a visual comparison of PAA-ED score and similarity between modelled and measured demand profiles, described in detail in [12], a PAA-ED score of less than 2.5 was determined to show high similarity between two distributions, a score between 2.5 and 3.5 good similarity, a score between 3.5 and 4.5 some similarity, and a score greater than 4.5 showed low similarity.

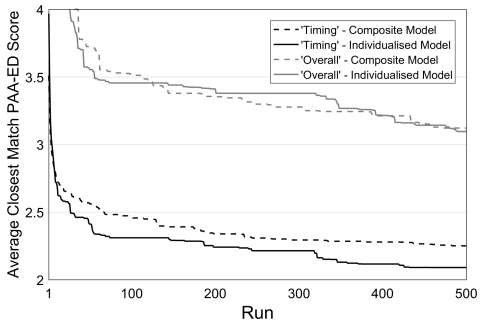

(a) HES dataset

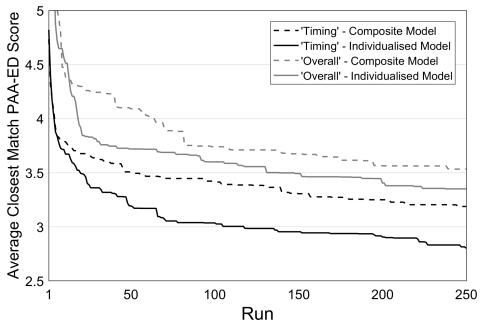

(b) 'Richardson'

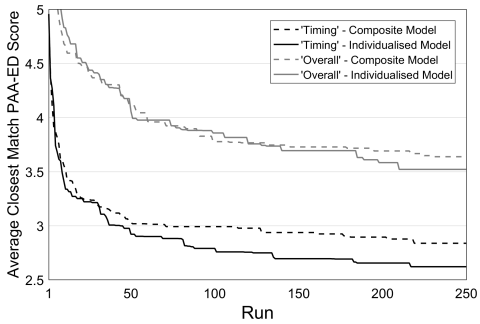

(c) REFIT

Fig. 15. Cumulative closest match PAA-ED score average per run comparison between 'composite' and 'individualised' models for different dataset-equivalent electricity demand models.

Five hundred Household Electrical Survey (HES) dataset household equivalent model runs (based on known household characteristics) for both 'composite' and 'individualised' methods were compared. After each run and for each of the modelled HES-equivalent households, the cumulative lowest PAA-ED value for all runs completed is determined between each modelled and actual HES household and the average calculated. This average PAA-ED value is a simple measure of the ability of the model to generate demand profiles consistent with the range of behaviours observed in the measured data and the progression of average results with runs performed gives an indication of the speed with which each highly probabilistic method identifies real and representative patterns of behaviour.

The results for the full 250-household HES-equivalent model are inconclusive, with similar results for both methods. However, for the 26 HES households which were monitored for

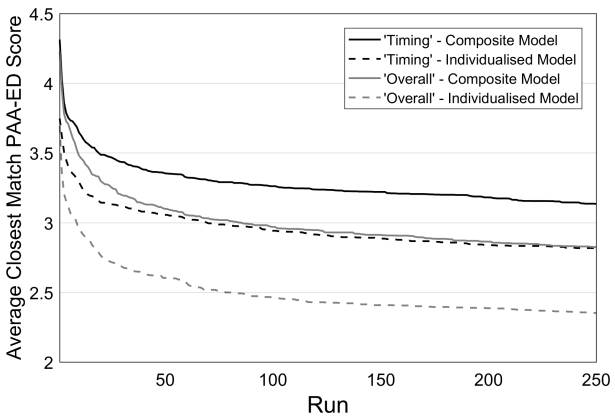

Fig. 16. Cumulative closest match PAA-ED score average per run comparison between 'composite' and 'individualised' models for EST dataset-equivalent hot water demand model.

longer than 28 days (between 61 and 249 days with an average of 125 days), the average for both methods is significantly lower, and there is an improvement in the 'Timing' value after 500 runs for the 'individualised' model with an average of 2.09 compared to 2.25 for the 'composite' model (see Figure 15(a)). This duration-dependent performance suggests that the length of the analysis period is also important, with 28-day profiles being significantly more erratic than longer duration profiles and therefore more difficult to replicate.

Table 3
Final PAA-ED score for comparison of annual duration datasets with the equivalent model output based on best match per modelled household after 500 runs.

	'Composite'	'Individualised'
Richardson - Timing	3.19	2.80
REFIT - Timing	2.84	2.62
Richardson - Overall	3.53	3.34
REFIT - Overall	3.64	3.51

Similar PAA-ED analysis with the Richardson [4] and REFIT [29] electrical demand datasets (20 and 21 households respectively), and equivalent generated model outputs, which are all of a 1-year duration, show a clear performance improvement for the 'individualised' method basis (see Figure 15(b) and (c)) with the final 'best-match' results as shown in Table 3 showing an improvement for all cases.

5.3.2. Hot Water Demand

PAA-ED analysis was also undertaken for the hot water demand sub-model comparing the Energy Savings Trust (EST) dataset used for sub-model calibration with the dataset-equivalent modelled population.

For both the 'Timing' and 'Overall' assessments (see Figure 16), the 'individualised' method shows significantly better performance than the 'composite' method. The improvement is also greater and more consistent than for the electricity demand sub-model. The similarity assessment of individual results also shows a distinct improvement in the number of 'High' (PAA-ED score < 2.5) and 'Low' (PAA-ED > 4.5) similarity results as shown in Table 4.

Table 4
Hot water model closest cumulative match similarity analysis range results for the 'composite' and 'individualised' EST household equivalent models after 250 model runs. (High = PAA-ED score < 2.5, 'Good = 2.5-3.5, 'Some' = 3.5-4.5, 'Low' = 4.5+)

Similarity	Avg.	'High'	'Good'	'Some'	'Low'
'Timing' - 'Individualised'	3.5	62 (62%)	31 (31%)	7 (7%)	0 (0%)
'Timing' - 'Composite'	3.5	46 (46%)	33 (33%)	12 (12%)	9 (9%)
'Overall' - 'Individualised'	3.5	37 (37%)	40 (40%)	21 (21%)	2 (2%)
'Overall' - 'Composite'	3.5	28 (28%)	41 (41%)	21 (21%)	10 (10%)

5.3.3. Similarity Analysis Conclusion

The results of the demand model analysis for the developed individual-calibrated ('individualised') method show that it performs better than the group-calibrated ('composite') method, particularly where annual data is available for comparison. The slightly weaker replication of the group-average occupancy characteristics with the addition of the individual occupancy transition timing adjustments shown in 5.1.2 is therefore outweighed by the improvement in capturing individual demand behaviours. Moreover, the main source of the current weaker occupancy replication performance, the sleep-transition timing element, can be improved within the 'individualised' occupancy method with improved calibration.

The significant performance improvement seen for the hot water demand sub-model, with the significant model performance improvement shown in Table 4, suggests that the impact of the individualisation method can be significant. Applied to the 'High' and 'Very High' cycle volume ranges, the 'individualised' method impacts on 63% of hot water use with a single behaviour adjustment per household.

The electricity sub-model analysis is more complex, incorporating multiple appliances with different, and potentially conflicting, behaviours and power profiles, which account for a far smaller proportion of overall demand, and were monitored for shorter periods. It is therefore more difficult to compare overall electrical demand model output with existing small-scale datasets. However, the improvement in individual appliance behaviour replication and relatively low correlation in behaviour between appliance use within households, and the smaller, but meaningful, improvement in overall demand profile replication, suggests that the new methods have merit where individual behaviour replication is required. The output is also potentially useful for analysis of demand shifting potential for specific appliances and demands.

6. Discussion

The presented modifications to the previously developed occupancy and energy demand model [12], have been demonstrated to both reduce unrepresentative convergence to average

behaviours and better generate more diverse individual occupant and household behaviours.

It is acknowledged that in the further manipulation of the directly generated calibration data, an element of potential inaccuracy is introduced to the process. However, the selected manipulations were chosen specifically to ensure that any introduced variation was centred on the mean of the overall population and allowed to vary only within controlled limits constrained by realistic variance from average behaviour. Over a significant number of runs it was confirmed that the results finally converge back to the calibration dataset basis, maintaining statistical consistency. Individual profiles were also scrutinised for any evidence of unrealistic outcomes. The main conclusion was that the model retains an inability to capture the more extreme behaviours seen in the measured data, which are potentially outwith the capability of a probabilistic model, rather than generating results that outlie measured data.

While the 'individualised' model basis provides an output that is more accurate for individual households, the 'composite' model basis retains advantages for specific occupancy and demand problems. The 'individualised' model requires a significant number of runs to provide a representative sample. Conversely, whilst the 'composite' model is unrepresentative of individual households, it will capture average behaviours within a small output sample. The preferred method will therefore depend on the scale and type of analysis, and the need to balance capturing individual-specific behaviours with overall behavioural consistency. Further work is required to determine appropriate methods to filter multiple 'individualised' runs to generate representative samples for use in wider building simulation analysis or whether the degree of individualisation can be appropriately manipulated to reduce the number of individual simulations required.

7. Conclusion

As identified by the recent IEA Annex 66 activity, integrating occupant behaviour is critical for state-of-the-art building simulation. As shown in this paper, existing calibration approaches, using composite data from multiple households, does not generate output that fully reflects individual behaviour variation for occupancy, individual appliance use or overall demand.

In the absence of significantly larger and more integrated occupancy and demand datasets, it was necessary to modify existing modelling methods using identified underlying statistical relationships inherent in the calibration data in combination with analysed variations in average behaviour. This ensured that the available data is used to its fullest extent, that the degree of convergence in the results was significantly reduced and to a level that ensured annual model output retains realistic variance and the overall statistical characteristics of the calibration population.

The next step is to integrate the developed occupancy and demand model output with building simulation to ensure that the impact on heating and cooling use both directly in terms of use timing and also indirectly via occupant and equipment gains. This can be used to stress-test fabric and HVAC system options under a realistic range of potential scenarios rather than relying on average behaviour or potentially unrepresentative archetype models.

Acknowledgements

The authors gratefully acknowledge the financial support received for this work from the BRE Trust.

Open-access Model

The occupancy and demand models detailed in this paper (named OccDem), both the 'composite' and 'individualised' models via a user-selection, are available from *http://www.esru.strath.ac.uk/applications/occdem/index.htm* via the most recent upload.

References

[1] R. Yao, K. Steemers, A method of formulating energy load profile for domestic buildings in the UK., Energy and Buildings 37 (6) (2005) 663–671.

[2] Y. G. Yohanis, J. D. Mondol, A. Wright, B. Norton, Real-life energy use in the UK: How occupancy and dwelling characteristics affect domestic electricity use., Energy and Buildings 40 (6) (2008) 1053–1059.

[3] C. Walker, J. Pokoski, Residential load shape modelling based on customer behavior., Power Apparatus and Systems, IEEE Transactions on (7) (1985) 1703–1711.

[4] I. Richardson, M. Thomson, D. Infield, C. Clifford, Domestic electricity use: A high-resolution energy demand model., Energy and Buildings 42 (10) (2010) 1878–1887.

[5] J. Widén, E. Wäckelgård, A high-resolution stochastic model of domestic activity patterns and electricity demand., Applied Energy 87 (6) (2010) 1880–1892.

[6] U. Wilke, Probabilistic bottom-up modelling of occupancy and activities to predict electricity demand in residential buildings., Ph.D. thesis, École Polytechnique Fédérale de Lausanne (2013).

[7] E. McKenna, M. Thomson, High-resolution stochastic integrated thermal–electrical domestic demand model., Applied Energy 165 (2016) 445–461.

[8] R. Baetens, D. Saelens, Modelling uncertainty in district energy simulations by stochastic residential occupant behaviour, Journal of Building Performance Simulation 9 (4) (2016) 431–447.

[9] D. Fischer, T. Wolf, J. Scherer, B. Wille-Haussmann, A stochastic bottom-up model for space heating and domestic hot water load profiles for german households, Energy and Buildings 124 (2016) 120–128.

[10] A. Taniguchi, T. Inoue, M. Otsuki, Y. Yamaguchi, Y. Shimoda, A. Takami, K. Hanaoka, Estimation of the contribution of the residential sector to summer peak demand reduction in japan using an energy end-use simulation model, Energy and Buildings 112 (2016) 80–92.

[11] G. Flett, N. Kelly, An occupant-differentiated, higher-order Markov Chain method for prediction of domestic occupancy., Energy and Buildings 125 (2016) 219–230, http://dx.doi.org/10.1016/j.enbuild.2016.05.015.

[12] G. Flett, N. Kelly, A disaggregated, probabilistic, high resolution method for assessment of domestic occupancy and electrical demand, Energy and Buildings 140 (2017) 171–187.

[13] G. Flett, Modelling and analysis of energy demand variation and uncertainty in small-scale domestic energy systems, Ph.D. thesis, University of Strathclyde (2017).

[14] K. B. Janda, Buildings don't use energy: people do, Architectural science review 54 (1) (2011) 15–22.

[15] N. Saldanha, I. Beausoleil-Morrison, Measured end-use electric load profiles for 12 canadian houses at high temporal resolution, Energy and Buildings 49 (2012) 519–530.

[16] D. Murray, L. Stankovic, V. Stankovic, An electrical load measurements dataset of united kingdom households from a two-year longitudinal study, Scientific data 4 (2017) 160122.

[17] A. Tversky, D. Kahneman, Belief in the law of small numbers., Psychological bulletin 76 (2) (1971) 105.

[18] I. Richardson, M. Thomson, D. Infield, A high-resolution domestic building occupancy model for energy demand simulations., Energy and Buildings 40 (8) (2008) 1560–1566.

[19] M. Muratori, M. C. Roberts, R. Sioshansi, V. Marano, G. Rizzoni, A highly resolved modeling technique to simulate residential power demand., Applied Energy 107 (2013) 465–473.

[20] H. Meidani, R. Ghanem, Multiscale Markov models with random transitions for energy demand management., Energy and Buildings 61 (2013) 267–274.

[21] A. J. Collin, G. Tsagarakis, A. E. Kiprakis, S. McLaughlin, Development of low-voltage load models for the residential load sector., Power Systems, IEEE Transactions on 29 (5) (2014) 2180–2188.

[22] D. Aerts, J. Minnen, I. Glorieux, I. Wouters, F. Descamps, A method for the identification and modelling of realistic domestic occupancy sequences for building energy demand simulations and peer comparison., Building and Environment 75 (2014) 67–78.

[23] M. Kleinebrahm, J. Torriti, R. McKenna, A. Ardone, W. Fichtner, Using neural networks to model long-term dependencies in occupancy behavior, Energy and Buildings 240 (2021) 110879.

[24] A. Grandjean, J. Adnot, G. Binet, A review and an analysis of the residential electric load curve models., Renewable and Sustainable Energy Reviews 16 (9) (2012) 6539–6565.

[25] Sociaal en Cultureel Planbureau, Tijdsbestedingsonderzoek 2005 - TBO 2005. (2005).
URL `http://dx.doi.org/10.17026/dans-znn-5xvz`

[26] Office of National Statistics, United Kingdom Time Use Survey, 2000., [data collection]. 3rd Edition. UK Data Service. SN:4504. (2003).
URL `http://dx.doi.org/10.5255/UKDA-SN-4504-1`

[27] DECC, Household Electricity Survey., [Accessed Nov. 26, 2013] (2012).

[28] Energy Savings Trust, Measurement of Domestic Hot Water Consumption in Dwellings. (2008).
URL `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/48188/3147-measure-domestic-hot-water-consump.pdf`

[29] D. Murray, J. Liao, L. Stankovic, V. Stankovic, R. Hauxwell-Baldwin, C. Wilson, M. Coleman, T. Kane, S. Firth, A data management platform for personalised real-time energy feedback., in: Proceedings of the 8th International Conference on Energy Efficiency in Domestic Appliances and Lighting, 2015.