Centre for
Sustainable Energy

An analysis of D3 in DECC's energy system models

Report to Department of Energy & Climate Change

Tom Hinton and Joshua Thumim

July 2014

Centre for Sustainable Energy

3 St Peter's Court
Bedminster Parade
Bristol BS3 4AQ

0117 934 1400
www.cse.org.uk
info@cse.org.uk

We are a national charity that shares our knowledge and experience to help people change the way they think and act on energy.

(Reg charity 298740)

Sign up to our monthly enews at www.cse.org.uk or follow us on Twitter at @cse_bristol

Contents

1 What is this document?

The Department for Energy and Climate Change (DECC) have retained us, the Centre For Sustainable Energy (CSE), to review their approach to modelling *D3*, the family of technologies and interventions which affect the use and production of energy on the demand side of the energy system. In this document, we hope to present a picture of:

1. What challenges D3 poses for energy modelling
2. How DECC's models represent D3
3. How the models are connected together
4. How the models are already used to consider D3 questions
5. Where the models may fail to capture D3 costs or benefits, either because of the way they are used, or due to fundamental limitations in the design
6. What DECC could do to improve their models or use of models in the analysis of demand side policies and D3.

1.1 What is D3?

As mentioned, D3 is concerned with the demand side - interventions whose effects are meeting demands themselves, say by reducing the energy required to deliver the service, or changes which meet demand by moving energy around near the ends of the distribution network, through microgeneration, distributed storage and so on, fall into the D3 category.

D3 is named for a division of these things into the three categories of *demand reduction*, *demand response*, and *distributed energy*.

Demand reduction measures are those which reduce overall demand, either by increasing efficiency (so the same useful work is done for less energy input), or by reducing the desire for useful work. Examples for the first category would include insulation, better heating controls[1], or more efficient machines. An advertising campaign which persuades everyone to wear more jumpers and turn down their thermostat would belong to the second category.

Demand-side response (DSR) technologies are those which allow the supply side to automatically control demand in some way, such as with real-time electricity pricing, or by providing demand destruction to the grid as an option for meeting supply (sometimes referred to as *negawatts*).

Demand response is typically thought of as a mechanism for reducing *peak* loads by moving some of the peak to times of day when demand is lower. Consequently it may not reduce overall energy demand, but instead affects the peak power demand.

Distributed energy encompasses technologies which produce heat or electrical power near the end of the distribution network, and energy storage systems which can balance loads at a similarly local scale[2]. Examples would include solar photovoltaics or hot water, small-scale combined heat and power, heat storage, or small (say ~2kW) wind turbines.

This taxonomy is not of critical importance; particular technologies within each category can be quite different (for example insulation, efficient boilers and woolly jumpers all provide demand reduction), and may arguably belong to more than one category (demand response and energy storage can each be thought of as a form of the other, for example).

1.2 Method

The modelling approach to D3 has two key parts: firstly, the models themselves, which by nature have strong representations of some parts of the world and weaker representations of others, and secondly the organisational context which causes them to be used one way or another.

To try and understand these, we have (where possible) looked into the implementation details of the models, and spoken with their users about their capabilities and limitations, and the ways in which they are normally used. In doing this we have tried to capture the following:

- For each model, how is the demand side represented, in particular:
    - Which demand side effects are presented as inputs, and how are the input values produced?
    - What demand side effects are predicted as outputs, and where are the outputs used?
    - When the model is simulating part of the demand side, how does that simulation work?
- For models and their combinations:
    - What normally drives their use?

---

[1] Heating controls increase the efficiency of a heating system in terms of the amount of energy required to deliver the desired result, i.e. a comfortable temperature at particular times of day. They could be thought of as reducing demand (they produce a demand at the boiler, after all), but here we consider them as part of the machinery that meets the real human need for useful work.

[2] Electricity storage could reasonably also be considered as a form of demand response, shifting grid load from times when electricity is plentiful to times when it is scarce (whereas demand response shifts demand from times of scarcity to times of plenty).

  - Who are they typically run for, how often, and under what conditions (changes to inputs, etc.)?

The inter-model connections naturally lead to an iterative investigative process; one model is joined to another, and so we have had several rounds of dependency chasing. Since the modelling network and constituent models are quite large, and involve some work whose provenance is external to DECC, we may not have formed a comprehensive picture, and our choice of starting point may have biased our choice of models and policies to investigate in a less than ideal way. That said, we have developed a reasonably detailed picture of the models, their connections, and their uses, which is presented in this report.

Our initial objective was to consider some *partial system models*, which we imagined to be developed for answering specific policy questions, and some *whole system models*, which are those that attempt to model close to all of the energy system, and we expected to be used for the production of consolidated figures and the assessment of systemic effects. Some of these models are *optimising* models, which attempt to determine the *best* solution which satisfies some constraints, and some are *simulation* models, which instead are designed to produce a plausible future scenario under some assumptions about how a future state can be derived from a past one. Table 1 names and classifies all the models referred to in this document.

We would like to thank all of the analysts we have spoken with at DECC and Department for Communities and Local Government (DCLG), who have been courteous and helpful to a fault, and very generous with their time; these are:

- Simon Green (Dynamic Dispatch Model (DDM))
- Marjorie Roome (Energy Demand Model (EDM)/Unified Emissions Projection (UEP))
- Bevan Freake (Energy Systems Modelling Environment (ESME))
- Jonathan Roberts (ESME)
- Roger Lampert (Combined Heat and Power (CHP) Model)
- Ting Ho (Renewable Heat Incentive (RHI))
- Eamon Graham (Transform)
- Paul Decort (Building Regulations)
- Guido Coco (Smart Meters)

1.3 What else is in this document?

The most detailed descriptions of each model are given in section 6. Section 2 explains some of the difficulties in modelling the demand side, section 3 gives an overview of where the demand side has been represented in the models we have found out about, section 4 gives some suggestions on how DECC's existing tools could be put into service for modelling D3, and section 5 highlights those limitations which seem more problematic and may require some work to address.

1.4 What is not in this document?

The negative space is also important here; there are several things this document does not try to provide:

1. A review of the effect of existing policy decisions on D3, so we have not undertaken to find out whether the specifics of existing targets or regulations promote or counter the adoption of D3 measures, or the development of policies targeting them.
2. *Quantitative* determinations of missed costs or benefits due to the approach taken so far; if existing modelling work has been unable to capture some aspects of D3, it is because it is difficult to model by taking a high-level overview, which is all we could admit within the scope of this work.

Table 1: Type is **W**hole system or **P**artial system and **O**ptimising or **S**imulating, and Source code and Owner whether or not we have seen the source or spoken with the author or an expert user respectively. The UEP and DDM are starred because they are primarily simulations, but use small optimisations in solving the relationship between price and demand, and the choice of new plant to construct when forecasting investment decisions.

Model	Type	Source	Owner	Notes
MARKAL	W/O	N	N	We have not examined MARKAL; it is mentioned only because it has been used to forecast demand past 2030 for the EDM.
ESME	W/O	Y	Y	
DDM	P/S*	Y	Y	
GDHM	P/S			
FITS	P/S			Due to time pressures faced associated with the delivery of the Solar Strategy we were not able to meet FITs model analyst; what we have been able to say in this document is based on published material from IAs, and our recollections of a previous version of the model which we decided not to bid for work on.
RHI	P/S	N	Y	
EDM	P/S	Y	Y	The EDM is partial in the sense that it doesn't model the supply-side for electricity.
UEP	W/S*	Y	Y	In this document, when we refer to the UEP we are describing a running combination of the DDM and the EDM.
DIMPSA	P/S	Y	Y	We have spoken with the authors of DIMPSA here at CSE for background understanding of how it incorporates different effects, but we have not written much about it here as it seems only to be used for the calculation of distributional impacts in the prices & bills report.
Building regs.	P/S	N	Y	We have spoken with DCLG about how Building Regulations are modelled, but we have not looked at any actual model implementation details.
Products	P/S	Y	Y	We have had access to a single instance of a products policy model, which is hopefully representative.
CHP	P/S	Y	Y	Although Roger Lampert has very kindly taken the time to produce a version of the CHP model which included no sensitive information, we have naturally not been able to see the real thing because it contains commercially sensitive material (some of which is not visible to DECC in any case).
Smart Grids	P/S	N	Y	We have not been able to see the smart meters model due to time constraints, but we have spoken with some of its current owners.
NHM	P/S	Y	N	The NHM is still being rolled out for actual policy modelling work; however, having developed it for DECC, we feel we are able to give a fair view on its strengths and weaknesses.

1.5 Summary findings

D3 is potentially a very difficult area to model in several ways, which we cover in (2); there are lots of possible interactions, and the potential significance of each is not obvious. Our review of DECC's models suggests that although each individual modelling effort is done carefully and to a good standard, the integration of results from multiple models and the influence of one model onto another is less well done. Although DECC's analysts make considerable effort to account for possible overlaps between their models, the difficulty of getting this right increases with each model added to the mix and each type of interaction that is allowed. This effect is exacerbated by the use of tools which are not designed for programming complex software systems with a lot of interacting parts (particularly Excel), the frequent rotation of staff and pressure to produce answers at speed, and the fact that a lot of important programming work has been outsourced. These factors have combined to produce a situation where the understanding of how things fit together is quite fragmented, and where exploring systemic effects is very difficult. We also note a lack of modelling around heat networks, which seem like they may be particularly difficult to analyse, but could also be potentially significant D3 interventions. This said, we believe there are several steps DECC could take to improve the situation.

1.5.1 Recommendations

DECC should:

1. Implement some processes and form a team of people whose job is to keep track of how their portfolio of models are connected up, what factors they address, and where there may be potential issues synchronising their side-effects on each other.
2. Undertake to quantify the significance of the many possible D3 effects that could be important. Some of these may be simulated using models like Transform or the DDM; others will require some novel research. Some specific suggestions to support this are:
    - Model distributions of inputs where the phase space is large/complex enough to permit a very wide range of possible outcomes (see (2.2.2) and (4.4)).
    - Generate a table of DDM whole-system cost results from a set of different absolute demand and peak:base ratio input combinations (see (4.2)).
    - Using Transform in a similar manner, generate a table of network costs and benefits from a set of combinations of the inputs to which the model is found to be sensitive (again, see (4.2)).
    - Develop a spatial model of heat networks and their relationship with power networks, because of the particular spatial characteristics described in (2.2.2).
3. Adopt a different set of standard tools for any modelling work which is not of the simple, exploratory sort which Excel is suited to - the lack of any general-purpose programming languages really stands out here. This may also imply training people in how to do analysis programatically, and trying to recruit people with more experience in programming or data science, although many of DECC's staff already have some programming experience. More detail about this is given in (5.5); the key points are:
    - Adopt automation and integration tools to manage the interactions between different models.
    - Introduce comprehensive version control for all models and model runs.
    - Try and improve the IT restrictions which prevent the widespread use of tools other than Excel.
4. Consider designing its models for interaction; working out the details of this are beyond the scope of this document, but would involve some kind of centralised mechanism for interleaving model runs in time and allowing models to operate on a shared representation of the world. We suggest the following steps in this direction:

    - Develop formal interfaces for Transform and the DDM to allow consistent modelling of generation and distribution costs using the best available tools, as suggested in (4.4). We would suggest exposing these interfaces for consumption from other tools, including Excel, and investing in the infrastructure needed to let analysts invoke these models without having to ask their DECC owners for time.
    - Ensure that policies which impact on generation and distribution costs (e.g. the Smart Meters model, which assumes peak load shifting) conform to this interface, and provide inputs/consume outputs from Transform and the DDM. An example use for this would be: policy models generate a distribution of demand changes over electrical substation types, whose combinations Transform consumes to determine reinforcement costs and to produce a grid-level load profile as an input to the DDM.
    - Develop a central representation of the world which is shared across models via standardised interfaces. This would require documentation of all interfaces as a subtask which would have value in its own right.
    - Use this central state representation to improve modelling of causal relationships by interleaving events in simulated time; the reasoning behind this is explained in (5.4).

   Connecting DECC's existing set of models would be an expensive and difficult undertaking, so it might be wise to start with a smaller proof-of-concept effort to create a simple shared-state simulator which interconnects some lightweight models before looking at refitting older models.

5. Where possible, try and disentangle the economic and behavioural parts of its models from the physical modelling of the energy system.
    - Explicitly recognise the difference between physical and social modelling; a better integrated physical model of the energy system could then be controlled by different social models to investigate where they diverge, for example.
    - Develop a better physical energy model for housing - although the Standard Assessment Procedure (SAP) has a physical basis, it has a variety of behavioural and climatic assumptions "baked in", and is used to serve the two incompatible purposes of (a) providing building ratings which are comparable in the face of varying external factors (standard occupancy and mean climate, for example), and (b) predicting actual energy consumption for a specific house. In particular, in (2.2.1) we outline some of the ways this adversely affects modelling of heat pumps, because of their sensitivity to climate.

1.5.2 Example scenarios

These are some illustrative example scenarios taken from DECC's "Whole Energy Systems Questions". For each question we have considered which of DECC's models represent the important technologies and whether there are any clear gaps that are unaccounted for. We have attempted to choose questions for which a heavily analytic answer seems appropriate; questions about political feasibility are likely to be out of scope for most of DECC's models.

What is the impact of behaviour change on energy demand? The energy effect of specific behaviour change could addressed to some extent using *Transform*, the *National Household Model (NHM)*, and the *DDM*; the NHM provides some facilities for predicting the effect of changing domestic *heating* behaviour, and Transform and the DDM can together estimate some whole-system costs tied to differences in demand profile. However, predicting which behaviour changes are likely to happen is an entirely different matter.

How do we develop locally appropriate heat strategies? The National Heat Map provides some basic information which could be useful when thinking about localised heat strategies. DECC has no other models (that we know of) which operate at a local level for heat.

What happens to electricity peak with under/over performance of housing insulation policy? This question could be answered in part using the *DDM*, *NHM*, *ESME* and *Transform*. The NHM does not have low-level temporal resolution, and so cannot correctly answer questions about using heat storage, but it does have a detailed model of the housing stock. ESME can represent time-shifted demand, but has a limited view of the stock and a focus on optimisation rather than simulation. Transform and the DDM could be used to assess the system costs associated with particular insulation scenarios.

What will it take to get the housing stock up to scratch? This could be answered in part using the *NHM*, by asking what the minimum cost option to bring every house in the stock up to some standard is, but depends on good-quality data about the costs and effects of all the different technologies being considered.

How do technologies interact? This question is (a) of quite a wide scope, and (b) not well handled by many of DECC's models. It may require primary research.

What is the potential (and cost) to increase energy efficiency of buildings by 2050? The *NHM* may be able to answer this question to some extent for the domestic stock. The lack of a good data set for the non-domestic stock makes it a much harder question on that side.

Can we balance demand and supply of electricity and heat by storing energy for days/seasons? *ESME* is the only model DECC has with a good understanding of heat or power storage; however, considering (for example) maximal insulation of the housing stock together with electrified heating *and* the use of large heat stores would be difficult with ESME because of its simplified view of the stock.

2 What makes D3 hard to model?

There are two fundamentally confounding aspects to the demand side; firstly, it is not a single homogeneous object, but the sum of a lot of different points of demand, each under different conditions. Secondly, interventions may interact with each other at several different scales, in ways which depend on local conditions.

Alongside these two problems, there are practical problems about the availability of information. Some demand side technologies have only become possible quite recently, because of other enabling technologies, and so information about their consequences is quite limited. For example, predictions about the amount of peak load which could be shifted with widespread demand response can only be so well-founded, as demand response has not been rolled out on a national scale before. Similarly, the costs of mass-producing energy storage technologies are hard to forecast because the technologies are so new. That aside, let us consider the two more fundamental issues in some more detail.

2.1 Heterogeneity

D3 covers a range of interventions, each of whose suitability is dependent on a lot of different purely local factors. Disregarding any of the interactions between technologies, we suggest the following (incomplete) list of factors controlling the suitability of various D3 measures:

- The value of insulation depends on:
    - The preexisting demand for warmth (the less demand, the less demand can be saved)
    - The physical condition of the building to be insulated (easy/hard to treat, conservation areas, and so on).
    - To some extent, the remoteness of the location (it may cost more to insulate all the houses in the Hebrides than an equivalent set in the South East).

- Demand side response:
  - The feasible amount of demand response at a point of demand depends on the number of different demands and how large they each are - the cost of instrumenting and controlling a million 1 watt devices is likely to be much larger than instrumenting and controlling a single megawatt device.
- Heating technologies:
  - The feasibility and value of replacing simple boilers depends on the availability of fuel, on the heat load required, and the performance of whatever boiler already exists.
  - Heat pump performance is related to the factors affecting boilers, but also to the rise in temperature required between the hot side and the cold side; this in turn is controlled by some very local particulars like the normal ground temperature and the flow temperature in radiators.
- Microgeneration and storage:
  - Local planning constraints affect the practicality of some kinds of microgeneration.
  - Grid-connected devices require that the local grid connection be large enough to handle the flow to or from the device.

None of this list really addresses the human or policy factors which are also of great significance; to give a domestic example, it may be far easier to motivate people to improve their own homes than to motivate landlords to improve rental properties. Similarly some technologies are eligible for subsidies if the business or home purchasing them meets particular regional or socio-demographic criteria; this is to say nothing about the possible difficulties of understanding subsidies like the Energy Company Obligation (ECO), which are in principle controlled by all of the factors feeding into a SAP calculation. The fact that the demand side involves a large number of actors, many of whom are far from the “homo economicus” of classical economic theory, makes any economic modelling an exceedingly challenging task, quite apart from the direct physical factors affecting the technologies alone. This is in contrast to the hopefully more predictable behaviour of large corporate actors, who might be hoped to behave as economically rational agents because of their employment of many people working towards that goal. The question of where the ideal balance lies between effort spent attempting to model social processes (which may be intractable) and effort spent reducing the uncertainties and widening the scope of physical factors is a hard problem, and not one we feel equipped to answer.

2.2 Interactions at many scales

Accounting for non-additive effects is difficult; consider the following illustrative example:

- Loft insulation provides a benefit of 100 units; there are N_L un-insulated lofts.
- Cavity wall insulation provides a benefit of 100 units; there are N_C un-insulated cavities.

In this case, the available benefit is simply $100N_L + 100N_C$, which is easy to compute. Now consider the following confounding effect: loft insulation and cavity insulation *together* are of super-additive benefit, yielding an additional 100 units over and above the 200 you would expect. Without knowing how many houses have both an un-insulated loft and an un-insulated cavity, we can only bound the super-additive effect above and below. Although solving this only requires us to measure and think about one additional number, more realistic cases may be far worse, demanding many more additional dimensions be known about each category being considered. This has two key negative side effects; firstly, the set of possibilities to consider grows exponentially (what mathematicians refer to as the *curse of dimensionality*), and secondly the cost of acquiring detailed enough information about the real world grows with it. Coupled with de Moivre’s equation for sample power [10] this makes building a good model ever more expensive in computation and high-quality data as the number of nonseparable effects rises.

This trivial but hopefully illustrative example aside, let us consider some of the actual interactions on the demand side. As mentioned these can occur at several scales; those which are local to a particular building, or site, those which are relevant at the scale of the distribution network, and those which happen on the supply side in response to demand, but require some careful consideration. These kinds of interactions effectively compound all of the challenges posed by the heterogeneity of demand-side measures and the places where they might be applied. The effects described below are mostly of the simpler physical or technological kind, but the problem of many economic actors noted above recurs at these different scales of interaction as well.

2.2.1 Local interactions

These are interactions between demand side interventions which take effect after the connection to the distribution network, within a particular building or site.

Insulation and cold bridging Insulation typically has a super-additive benefit (as in the example above), because heat will take the path of least resistance to escape a building. Insulating one wall of a house will, to some extent, increase the heat flow out of the other walls, although not to such a degree as to make the insulation ineffective. Cold bridges provide a very strong form of this effect, because they have such high transmissivity that significant amounts of heat are lost through them.

This effect is ignored in simple models; for example, all SAP-type calculations produce a directly additive prediction, because the building's external surfaces are all presumed to be at heated temperature on one side and unheated on the other, and so are independent. In truth, insulating a wall has the effect of increasing the warm-side temperature of the other walls near it, which causes more heat to flow through them. Modelling the effect properly demands a realistic thermal model, likely using a finite element method, which would be intractable for many houses, but an estimate of its significance could be taken to produce some bounds for the error introduced by ignoring it.

Insulation and simple heaters Simple heaters here are those whose efficiency isn't changed much by the amount of power they have to supply; electric heaters and standard gas boilers are examples for these. When a heater is being installed alongside some insulation, the size of heater required is reduced, and so is the benefit produced by any efficiency increase. In this sense, insulation and boiler upgrades are antagonistic; each reduces the benefit associated with the other.

Insulation, radiators, the weather, and heat pumps Heat pumps[3] are slightly confounding; the performance of a heat pump is bounded above (according to Carnot's theorem) by $T_h/(T_h - T_c)$, where T_h and T_c are the hot and cold side temperatures. Rewriting as $1/(1 - T_c/T_h)$, we can see that as the cold side temperature falls below the hot side temperature, so too does the performance.

Three key factors control the difference between hot and cold side:

1. Required power output

   This is determined by how fast heat flows out of the building: the heat flow out of a building in a steady state must equal the heat flowing out of the radiators[4]; this in turn depends on the temperature difference between the radiators and the air, so that an increased power requirement demands a higher flow temperature[5], and so a lower coefficient of performance.

   Better insulation decreases the heat loss rate, and so increases the coefficient of performance because the radiators may be colder.

2. Radiator size

---

[3] Condensing boilers are also affected by similar factors, although not as drastically as heat pumps. The amount of heat that can be recovered in the condensing phase depends both on having a high enough return flow temperature and on the temperature and relative humidity of the surrounding air.

[4] Disregarding internal and solar gains, etc.

[5] For most radiators power output is proportional to something around $\Delta T^{1.3}$, with the constant being determined be the radiator area.

The other factor controlling radiator power output is area; increasing radiator area decreases the required flow temperature for a given output, and so increases the coefficient of performance.

3. Climate

   Finally, the *external* temperature is determined by such factors as the local climate, the thermal mass of the heat store the pump is connected to (if ground-source) and how much heat has entered the store, and the current weather (for air-source pumps). Because the performance function for heat pumps is typically not linear (so a cold day's losses may not be offset by a warm day's gains), simple mean temperatures are not sufficient to predict seasonal performance. Instead some notion of the distribution of temperatures at times of high heat demand is required.

The related trade-offs between heat pumps, insulation, and installing new radiators will have some cost-effectiveness tipping points, depending on the real characteristic function of the heat pump and the true thermal properties of the house. A real heat pump has a coefficient of performance far below the Carnot limit, and likely of very different shape, with an absolute maximum constrained by the machine's size.

Consequently, the evaluation of heat pumps cannot be done independently of an understanding of the existing condition and possibilities for installation of insulation (including the super-additive effects of insulation described above), the climatic conditions (including some distributional properties), and the costs of potentially fitting new radiators, pipes and so on.

Electrical efficiency and heat demand From the point of view of a thermostat, all the electrical appliances in a building are just poorly-controlled electrical space heaters. Anything which improves the energy efficiency of appliances thus also increases the demand for space heat in that building which must be met by the heating system. In this sense appliance efficiency improvements and insulation or fitting a new boiler are complementary, because the increase in appliance efficiency increases the cost/benefit for a more efficient heater.

Electrical demand and demand response Naturally, demand response mechanisms and the level of demand are connected in a complex way. The amount of demand response available in some location is determined by the nature of the loads which are connected; canonically, cooling systems are ideal candidates for demand response, whereas lighting and cooking are good examples of demands which are likely to be harder to shift. Consequently to relate demand reductions to demand response, we must understand which kinds of demand are being reduced, and how the means of their reduction affects their "shift-ability" in time.

2.2.2 Distribution network interactions

These interactions take effect at the level of the distribution network; there are several related effects which are important here:

Network reinforcement costs The electrical distribution network has costs mostly associated with the *peak* load on the wires. These *reinforcement costs* pertain at all levels of the distribution network, and so are connected with the energy use in many buildings at once; for example a substation at the lowest level might support a few hundred houses on several different feeders. The cost of reinforcing the substation depends on the peak size and duration for each meter, feeder, and for the substation as a whole. Between meter and Grid Supply Point (GSP) level, there are multiple layers to the transmission network; each layer may need reinforcement, depending on changes to the loads below it.

Demand reduction, changes in the fuel mix used to meet final needs, microgeneration, demand side response and energy storage may all combine together to affect the magnitude and the duration of the peak. Furthermore, reinforcement costs are non-linearly related to these terms, and so the reinforcement costs associated with a particular mix of technologies must be understood:

1. With sufficient temporal resolution to determine the time of peak.

2. With sufficient resolution on the factors affecting technologies *and their combinations* to understand the peak correctly.
3. With sufficient resolution over typical *ensembles* of houses to understand the effects from many meters on the substations connected to those meters.

This is a fine example of how evaluating combinations of D3 technologies is made difficult by the heterogeneity of the demand side.

Imagine for example a situation where heat pumps and electric cars are both under consideration, with the hope that the cars will provide stored energy to support the heat pumps at peak time. The winter peak power delivered in the form of gas is several times higher than the electrical peak, so the reinforcement cost could be significant - in the worst case, links could need reinforcing all the way up the distribution network to GSP level. However, if every house with a heat pump has an electric car, and the electric car can provide sufficient stored power, no reinforcement may be needed. In between these limits lie many different situations depending on where and how many electric cars and heat pumps are connected, and how much power is required for heat, and so on.

Sufficient demand-side generation may also create an opportunity for a more significant change to the topology of the distribution network, with more local cross-links and smaller or fewer hierarchical connections to the higher voltage transmission network. It seems likely this would depend on whether the distributed supply could be made reliable enough to permanently reduce the peak demand on the grid.

Geographical factors for heat networks The financial and technical viability of heat networks depends on the positions and magnitudes of heat demands, and the costs of creating a network to connect them (digging up a park is probably cheaper than digging up a road, for example). As a result, estimating the scope for district heat will depend on having spatial information about heat loads, heat sources and the ways in which they can be connected.

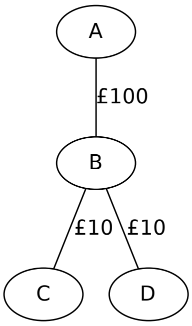

Consider the very simple example of the “city” on the right - each edge represents a place where we can feasibly put a pipe in the district heating system, labelled with the capital cost of installing the pipe, and each node is an address with a heat demand. Now let us say that a heat network will yield a certain amount of money per year from each connected address, depending on the demand for heat; in this case the economic decision to install a network is determined by whether the opportunity cost of capital investment is exceeded by the annual yield from the connected addresses. In the simple city on the right, say that A yields £ 50 per year, B £ 1 per year, and C and D each £ 20 per year, and that there is a base capital cost of £ 50 for any network. In this case, we have the following potential networks:

Network	Cost (£)	Yield (£/yr.)	ROI (%/yr.)
B, C, D	70	41	58
A, B, C, D	170	91	53
A, B, {C or D}	160	71	44
B, {C or D}	60	21	35
A, B	150	51	34

Now consider an otherwise identical city in which A has a yield of £ 20 per year, and C a yield of £ 50 per year - the table of possible networks becomes:

Network	Cost (£)	Yield (£/yr.)	ROI (%/yr.)
C, B, D	70	71	101
B, C	60	51	85
A, B, C, D	170	91	53
A, B, C	160	71	44
B, D	60	21	35
A, B, D	160	41	25
A, B	150	21	14

If we say that an acceptable ROI is 60%, then the first city has no feasible heat networks, whereas the second admits two; however, in most of their aggregate statistics the two cities appear identical. The distribution of connection costs and heat demands are the same, the number of addresses, and so on. It is only the topology which has changed, and yet the potential maximum yield has almost doubled! Although this is a simplistic example, it serves to highlight how much topology can affect the viability of heat networks; indeed the simplicity here only serves to underline how much worse the problem's complexity would be in the face of the many interacting factors and the much larger size of a realistic problem.

Interactions between heat and power Heat network economics are not only determined by the spatial distribution of heat loads and what needs digging up to connect them, but where the heat network is also providing electricity to the grid by the aforementioned network reinforcement costs! Consequently, to correctly assess a heat network supplied by a CHP generator, we must understand:

1. All of the factors mentioned above which affect network reinforcement costs, i.e. load profiles for different technologies, and the substations to which they are connected.
2. The heat demand profile (for heat-led networks), which will determine the co-generation profile, and so the effect on the distribution network.
3. All of the factors affecting heat network economics (including the topological effect described above).

Heat storage capacity As mentioned, heat pump performance is determined in part by the cold side temperature; this is the temperature of the heat store to which the pump is connected. The store temperature itself is affected by the amount of heat that has been put in, and the amount that has been taken out. For example, if a single house in a street uses a ground source heat pump it may perform differently than if every house in that street has one - the combined effect of many pumps may decrease the ground temperature sufficiently to affect performance.

This effect will depend on the thermal mass of the store, the heat input, and the rate of heat output required; these may make it irrelevant in most circumstances but it is not clear on inspection that it is negligible.

2.2.3 Generation effects

Having considered demands themselves and the network through which they are conveyed, we finally reach the effects on the supply side; in this we are talking about the supply of electricity at GSP level. As with the distribution network, the generation capacity required is determined by the annual peak; the cost of providing this peak capacity is amortised over electricity prices the for the rest of the time by some fairly complex market mechanisms. Furthermore, the cost of providing a certain peak is non-linearly related to the magnitude of the peak; this has two key consequences for modelling:

1. Avoided costs due to peak reduction cannot meaningfully be allocated to individual policies.
2. Competing policies or sets of policies cannot be selected by considering their marginal costs relative to business-as-usual.

To illustrate the second point, consider the following hyperbolically simple example: we have three policies, A, B and C, and twenty pounds to spend. Since A, B and C each costs the treasury £ 10 we can only choose two policies; for each policy we have produced an executive summary, as follows:

Policy	Cost	Avoided Gen.	Marginal Value	Air Quality	Total
A	10	100	10	2	12
B	10	50	5	1	6
C	10	30	3	4	7

Ranking the policies by their total benefit, we see that we should do policy A and policy C, but not B. *However*, the marginal value column has misled us; we have used a simple value of 10p per Wh avoided, but reducing generation by 140 Wh or more *in total* allows us to avoid building some plant, and so reduces the unit cost of electricity by some small amount. Under this model, the total value of A and B is £ 12 + £ 6 + £ X, which may well exceed the value of A and C with which we had naively compared it.

Obviously DECC does not engage in such simplistic and trivial exercises, but it *does* (possibly implicitly) trade off ensembles of policies against one another, and so ought to be considering the value of these avoided costs which belong to the holistic effect as well as the cost-benefit of each policy on its own.

3 Where is D3 in DECC's models?

More detail about all the models is in section 6, but first we present an overview of how most of the models are wired together, how they are used, and where there are representations of the demand side.

3.1 Overview

DECC has a lot of different models, which are sometimes used systematically for the production of certain key documents, and sometimes used in a more ad-hoc way to answer policy development questions or to provide some analysis to go into an impact assessment. The main modelling 'network' which is used for a systematic large-scale forecast is the Updated Emissions Projection model, or UEP. This is used to produce the eponymous document, which is a significant annual publication in its own right, and to generate values for the Interdepartmental Analysts Group (IAG) guidance on energy.

The UEP is centred on the EDM (6.1), which forecasts whole-system demand according to a counterfactual history evolving from 2000 in the absence of certain policies using some macroeconomic factors predicted by the treasury, ONS, OBR, and elsewhere. The EDM uses the DDM (6.2) to forecast electricity prices from demands, in a cycle - the price elasticity of electricity demand is represented in the EDM, and the cost of meeting that demand in the DDM. Along with the DDM, which produces prices to 2030, MARKAL (6.8) is used to produce a series of price forecasts extending from 2040 to 2050, and a linear interpolation is used in-between.

Along with this the UEP uses a set of time-series of *avoided consumption* by fuel and by sector from 2000 to 2035, each produced by one of many policy models; those in the past are backcasts of the effects particular policies have already had, and in the future are forecasts of the effects yet to come. It is not clear whether or by what means the backcasts are updated to reflect observation. Figure 1 shows the UEP as a diagram.

To understand the policy sub-models' effects on the EDM's outputs we can consider the size of the savings they predict - table 2 gives a lifetime summary for all the policy models which are connected to the EDM in the version we saw and figure 2 shows the change over time in the forecasts.

As the data shows, the lion's share of all forecast savings come from changes to Building Regulations, followed by the RHI, CERT policies (which are now completed), the Green Deal, and Products Policy. All of these policies involve D3 measures, and so deserve some consideration in this work.

Although the UEP is the largest frequently-run constellation of models, the picture above is not a complete view of how models are used within DECC. In particular, there are several other models which are not used for the UEP, but are used or will soon be in use for answering specific policy questions. Below is a figure showing the relationships that we encountered aside from (a) the UEP network described above and (b) the policy model to policy Impact Assessment (IA) relationship for each of the policy models; in the next section we give a summary of where different bits of D3 are in different models.

Table 2: Lifetime (2000-2035) policy savings entered into the EDM, broken down by policy. *Change* gives the absolute value (TWh) of all the changes in fuel consumption predicted, over all fuels and *Reduction* just the reduction in consumption of non-renewable fuels. Share is the percentage share in the total reduction. Most policies are expected simply to reduce overall consumption without shifting a lot of usage between fuels, but a few (particularly RHI and products policy) are changing the fuel mix as well as reducing non-renewable consumption. The reduction is plotted over time in figure 2. Values derived from the EDM's "Full Policy Savings" spreadsheet provided by Marjorie Roome (version: UEP48, v9, 26-Feb-2014).

Policy	Change (TWh)	Reduction (TWh)	Share (%)
Building regs 2002-2005	1075	1075	25.5
Building regs part L	611	605	14.3
RHI	1028	455	10.8
EEC1, EEC2, CERT	429	429	10.2
Green Deal	265	265	6.3
Products policy 2	277	250	5.9
Products policy 1	689	238	5.6
CERT 20% only + CERT Extension	185	185	4.4
Real time displays + Smart meters	158	158	3.7
Carbon Trust measures	148	148	3.5
CRC	116	116	2.8
Zero Carbon Homes	100	100	2.4
EPBD	77	77	1.8
Smart Metering	66	66	1.6
UK ETS	24	24	0.6
Warm front and fuel poverty	109	10	0.2
CESP	9	9	0.2
SME Loans	4	4	0.1
Salix loans	3	3	0.1
New CCAs	0	0	0

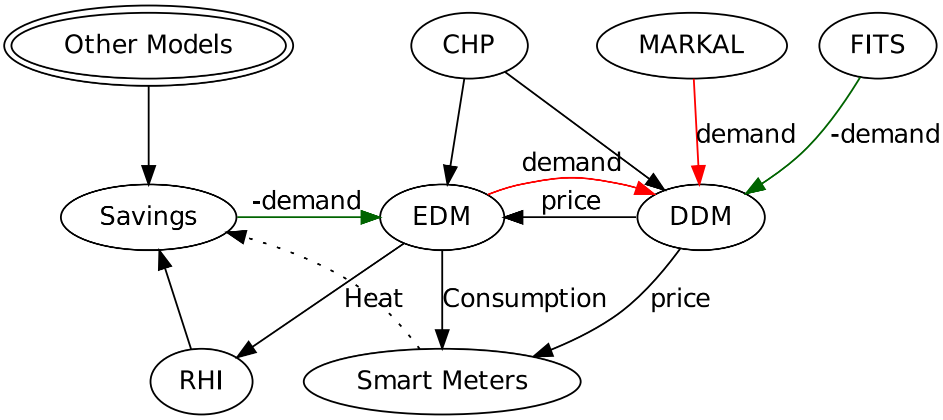

Figure 1: An overview of the modelling network used to produce most of DECC's whole-system outputs. Ovals are models, and arrows indicate input/output relationships. ESME and Transform are not shown in this model because they aren't typically connected to this network. "other models" node contains a set of per-policy models listed in table 2; RHI is broken out here because it is known to take a heat demand forecast from the EDM, and likewise the smart meters model takes base consumption.

3.2 Representations of D3

This subsection gives an overview of where different D3 measures and systems are represented in the models we have examined, summary information about the nature of the representation, and any potential limitations or missed interactions. The EDM is excluded from this list because although it must implicitly represent all of these factors in its econometric equations they are not present as separable parts.

Table 3: A summary table showing which technologies are in which of the key models. Y indicates a direct representation of some kind in a model, where a technology's effects are simulated in some way. I indicates an indirect representation, where a technology's effects may follow from a simulation in another model.

	ESME	BR	PP	GDHM	NHM	DDM	EDM	Transform	
Insulation	Y	Y		Y	Y	I	I		6
Heat Pumps	Y	?			Y	I	I		5
Boilers	Y	Y			Y				3
PV	Y					I		Y	3
District Heat	Y				Y	I	I		4
Solar Thermal	Y				Y		I		3
Supply Side	Y					Y		Y	3
EVs	Y							Y	2
Storage	Y					Y			2
Appliance Eff.			Y				I		2
Demand Response						Y	I	Y	3
Network Cost								Y	1
	9	3	1	1	5	7	6	5	8

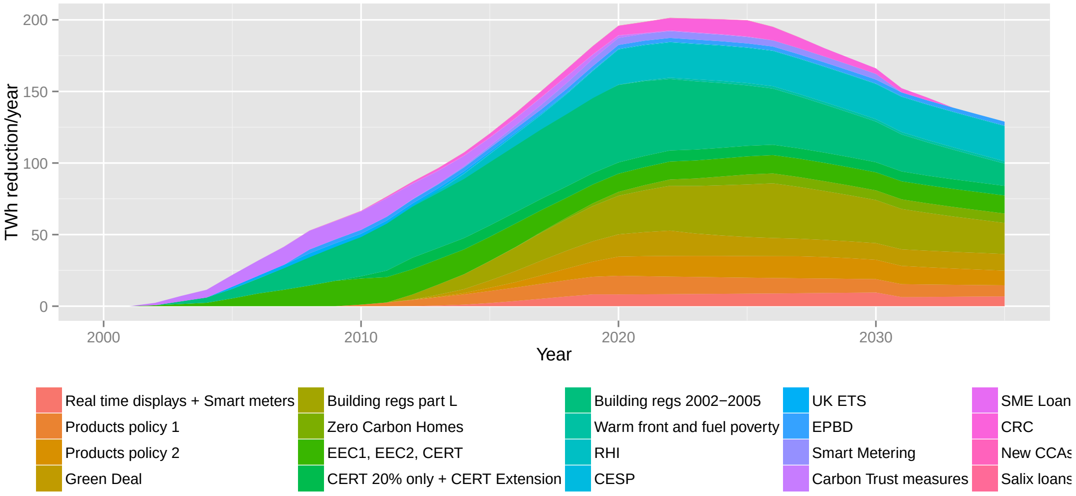

Figure 2: The reduction column from table 2 broken down over time. Note that the assumed savings decline after 2025 as the policies expire.

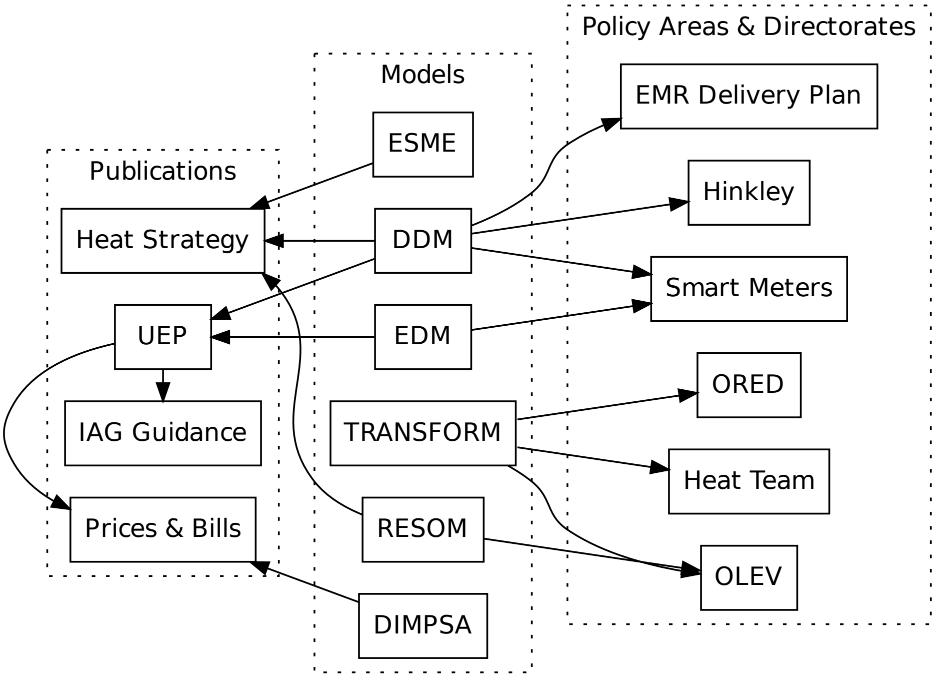

Figure 3: The uses of models for policy areas and preparation of documents which we have heard about, excluding the relationships between bespoke models and the policies for which they have been created (as these are self-explanatory). More detail about the uses can be found in the detailed model descriptions in (6).

3.2.1 Insulation

Building Regulations Since Building Regulations determine the thermal efficiency of new builds and some renovation works, the Building Regulations model ought to (and presumably does) incorporate some model of insulation and its uptake. However, we have not had access to the model or its authors and so can only speculate on how this works.

GDHM The Green Deal Household Model explicitly represents insulation as a measure that can be selected for installation by households; it has notions of internal and external solid wall insulation, cavity insulation and loft insulation. The energy savings the model uses are not calculated on-model, but by another tool we have not had access to; however, the algorithm is a SAP derivative, with the attendant limitations. It only represents the domestic sector.

NHM The National Household Model can be used to simulate insulation of most parts of a house (floors, walls, lofts, windows and doors), in accordance with a SAP-like energy calculation. The cost model can depend on the area insulated. Like the Green Deal Household Model (GDHM), it only represents the domestic sector.

ESME ESME represents space heat demand across multiple sectors, but is only able to change the thermal efficiency in the domestic sector, and with a simplistic representation which divides the housing stock into 12 different categories (three sizes crossed with four levels of thermal efficiency).

3.2.2 Heating systems

NHM The NHM can simulate the installation of heat pumps for domestic heat production; however, the simplest method from SAP is used, which just takes a single seasonal average coefficient of performance, with no consideration of emitter type, required heat output, or the temperature of the heat store.

The NHM can model changes in boiler efficiency for domestic heating. Its data for existing boilers are drawn from the Product Charateristics Database (PCDF) [7] using a fuzzy matching method on the boiler model information in the English Housing Survey (EHS), so insofar as the survey is representative this should provide quite a detailed picture of the existing situation.

New boilers can be installed in houses, with efficiency and cost depending on the likely size required and the effects of the change calculated using the SAP algorithm. SAP's model for conventional boilers is reasonably good (because they are easily characterised by one or two efficiency numbers), although the effect of regional humidity is not taken into account for condensing devices.

ESME ESME can simulate the use of heat pumps for space heat production, and can account for the influence of external temperature on heat pump performance to some extent, by using a different efficiency by season. The effects of regional temperature variation, required heat output and emitter size are not considered.

Simple boilers are also represented as devices which convert a fuel into space heat with a simple efficiency factor.

Building Regulations The Building Regulations model predicts significant reductions in gas consumption due to the eventual replacement of old low-efficiency boilers as they fail. However, as mentioned, the detail of this model is not something we have looked at in this project.

3.2.3 Appliance efficiency

Products policy The products policy models are concerned exclusively with appliance efficiency, and simulate the effects of reduced electricity consumption and the resultant increased demand for heat.

3.2.4 Energy storage

DDM The DDM represents grid-level pumped storage capacity. The pumped storage is not dispatched as a normal part of the merit order, but instead generates proportionally with demand except when it is being charged up. The model cannot currently simulate the purchase of additional energy storage.

ESME The most comprehensive handling of energy storage in DECC's models is in ESME; both heat and electrical energy storage across both daily and annual timescales can be represented and selected for construction.

3.2.5 Demand response

DDM Grid-level demand side response is represented in the DDM, although in a more restricted form than generation technologies. A predetermined price/capacity curve for each year of the run is defined in the model inputs, and included in the merit order when dispatching. However, the model does not have the freedom to suggest the outfitting of additional demand response as part of its construction forecast, and this cannot really be correctly handled as a special kind of plant with no consumption, because it should have a single capital cost but deliver different amounts of reduction at different marginal costs.

Also, the demand side response model that is included can represent the reduction of demand in response to the price signal, but not the shifting of demand away from peak time. This is because when DDM simulates the selection of demand response from the merit order, this does not lead to the simulation of increased demand at another point in time. In other words, the DDM's representation of demand response does not include modification of the shape of the demand profile, just its overall size.

Smart Meters The smart meters model does not simulate the intricacies of demand response directly; instead, it takes a demand figure from the UEP and some off-model analysis of the potential effect of smart metering in percentage terms, and multiplies the two together to determine the likely savings. It presumes that demand shifting will occur, but the shifted demand is not presented to other models which can understand it.

Transform Demand response is one of the technologies Transform can select as a way of resolving excess demand on the transmission network. We do not know how the amount of demand response that is possible has been modelled.

3.2.6 Distributed energy

FITS As mentioned, we have not reviewed the Feed-in tariffs model (FITS), and so have a very limited view of it. That said, it appears that most of its logic is involved in simulating the behaviour of investors, rather than the physical detail of where distributed energy technologies might go, how they could affect the distribution network, or how they might interact with smarter metering or energy efficiency measures.

ESME A raft of distributed energy technologies are present in ESME, but they all have the same basic form as a mechanism for converting one form of energy into another at some efficiency. The model has little spatial detail, and so the effects of distributed energy on the distribution system, and any other local interaction effects, are not represented here.

Transform Demand-side electricity generation is considered as one of the factors affecting network reinforcement costs in Transform; for example it can be used to determine the change in load profile caused by installing solar PV on some subset of substations.

3.2.7 CHP and heat networks

The CHP model The CHP model represents existing large-scale CHP installations (those above 50kW), and forecasts the development of analogous CHP developments under possible different future economic conditions. It does not model small-scale CHP or CHP heat networks.

EDM The EDM has some internal models for autogeneration, for sectors not covered by the CHP model. Those sectors that are covered by the CHP model are provided as inputs to the EDM. It uses these figures in the calculation of transmission losses and electricity demand.

ESME CHP is one of the technologies which ESME can use to meet energy service needs including heat, hot water and electricity; the model seems to select between heat and power led behaviour depending on which would be globally cost-optimal, which may not be typical of how CHP is actually used.

The model has technologies for district heat and hot water, but these do not use any spatial information in the calculation of cost or feasible capacity.

3.2.8 Generation costs

DDM The DDM is the best model DECC have for representing the cost of central generation. It simulates both electricity generation and the decommissioning and construction of new plant over time.

Transform Although Transform is designed for computing reinforcement costs, it includes a wider cost-benefit mode (as opposed to the narrow cost-benefit mode for network operators) which includes the costs from generation. This uses a dispatch model which looks similar to the DDM's, with a representative inventory of plant. Unlike the DDM, it does not do any investment modelling or forecasting of the change of plant over time.

ESME New plant can be constructed to meet demand in ESME; however, because ESME minimises the whole system costs, the decisions it suggests will be optimal for a central plan rather than a realistic forecast of investor behaviour. The formulation as a linear program will have omitted some properties which might be important, such as the fact that plant sizes are, in reality, quantized.

3.2.9 Distribution and reinforcement costs

Transform Transform is the most detailed model at DECC's disposal for predicting reinforcement costs. It has an understanding of the topology of the distribution network, the kinds of buildings which are attached to different substations, how various interventions at those buildings affect their load profiles and the different ways in which substations can be upgraded to meet increased loads.

RESOM We have not examined RESOM, but we understand that it has been used by Office for Low Emission Vehicles (OLEV) and the heat strategy team for assessing network reinforcement costs.

4 What more could DECC do to model D3?

This section contains some suggestions about investigations DECC could make using its existing set of tools to better assess the potential and challenge of D3 policies (the next section covers those places where we think DECC's models or techniques may be limiting the consideration of D3).

4.1 Quantifying the significance and potential of D3 interactions

Because this document is a review of *potential* gaps and opportunities, (2) is something of a laundry list of interactions, some of which may be insignificant. It may be possible to quantify their significance without having to develop fully representative models, or at least to produce upper and lower bounds on the magnitude of the various effects we have suggested. We would recommend that DECC convene a group of experts[6] to research this question, by the normal process of surveying whatever results already exist on the subject and then doing some exploratory thinking to fill in the gaps. Since the effects are systematic, and involve several different concerns this ought to be a group populated from across policy areas, and since the consequences may not be attributable to individual policies the work might have to be done under the *aegis* of the central modelling team or another equally central function.

4.2 Supply side costs

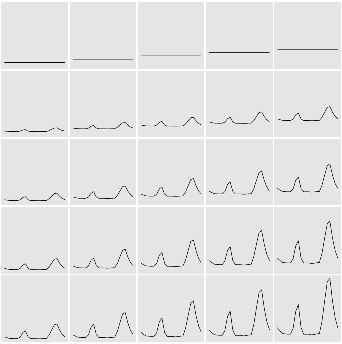

Figure 4: Load profiles with different absolute consumption (columns) and peak-to-base ratios (rows).

Although the generation cost has a complex relationship with the demand side, it is at least largely controlled by a relatively low-dimensional set of inputs: the load profile for the day containing the winter peak is very significant. We propose that DECC use the DDM (6.2) to produce a simple table giving a whole-system cost for every scenario drawn from a simple set characterised by two dimensions of the peak day:

1. An *avoided generation characteristic*, describing the value of avoiding different amounts of generation overall - this dimension is determined simply by the area under the curve.
2. A *peak/base ratio characteristic*, which summarises the value of flattening out the load profile - this is the ratio between the peak demand and the lowest trough.

These two variables can be adjusted quite simply; the first simply by multiplying through the load profile by some proportion, and the second using a transformation like $f'(t) = pf(t) + (1-p)\bar{f}$, where p adjusts the peak to trough ratio. This could produce a grid of load profiles, which could in turn be fed into the DDM to produce an easily understood summary of how the model responds, and hopefully how the real costs associated with generation vary with these two factors. The figure to the right is a simple illustration of this grid. Although most of the factors controlling the cost of a run are not being varied, a table of this sort would provide a simple view of how important these two factors are, and the benefit that might be available by carefully designing some complementary interventions.

---

[6] Preferably an *internal* group, as this would help to transmit and maintain the understanding of which interactions are important and need to be borne in mind within DECC, and possibly also help disseminate an understanding of who has modelled which parts of the system where.

4.3 Using Transform for reinforcement costs

Transform (6.9) is the most detailed model of the electricity distribution network readily available to DECC; we would recommend using it when making an assessment of the cost benefit of policies which will have effects on electricity demand.

Since peak load is the main determinant of reinforcement cost, the model could also be used to produce some figures on the potential costs of a range of exemplar scenarios much like those suggested for the central generation cost. Unfortunately this immediately runs up against the general problem of the demand side: any real intervention is going to vary against a range of dimensions which may be important (for example by region). However, a comprehensive investigation of Transform's behaviour might allow DECC to tease out the variables to which the model is most sensitive, and produce a look-up table cut along those lines which would be useful for other analysis; this may already have been done by its authors, or be scheduled within DECC already as a standard part of the model QA process.

4.4 Modelling ensembles

The previous two suggestions amount to the production of marginal cost/benefit tables, in the absence of policies or in a piecemeal per-policy way. Since the big difficulty with D3 comes in the varied interactions and non-additive costs at all levels of the system, it would ideally be better to consider ensembles of policies together. The two system costs which DECC is currently well equipped to consider are electrical generation and distribution costs, so we would suggest starting with these. To make modelling the shared effects of policies in these areas easier DECC should determine a common interface for each of Transform and the DDM, and ensure that policy-specific models produce outputs conforming with these interfaces; for example, policy models could generate a distribution of demand changes over electrical substation types, whose combinations Transform could consume to determine reinforcement costs *and* to produce a grid-level load profile for introduction into the DDM. Although this suggestion excludes feedbacks, in particular anything about the price elasticity of demand, it would be a good starting point. It is not too different to how the EDM, DDM and policy models are connected for the purposes of the UEP; the suggestion is simply to take this approach a step further, producing an automated continuous whole-system evaluation.

A simple diagram for the suggested system is shown below; the additional integration piece is shown as the dotted ellipse on the left, being fed by policy models A B and C. The system would ideally be automated, accepting and storing output scenarios from the input models and producing whole system costs for their combinations with little or no user interaction[7]; this automation would make considering holistic effects an easier habit to adopt, and produce a shared overview of the scenarios under consideration for different teams to look at.

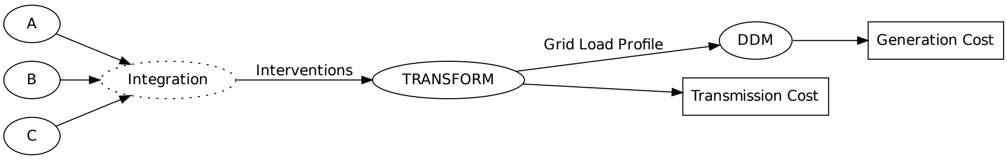

Some more thoughts on the potential for using some automation within DECC are described in (5.5). With some more attention to designing models for interconnection, in particular with respect to issues about causality (5.4), a system like this could support the flexible and interactive development style which DECC's modellers currently enjoy whilst enabling a much better holistic integration of the different scenarios that are being considered.

Since it may be difficult for all policy models to produce a good distribution of effect over substation types,

---

[7] This is a similar idea to *continuous integration*, a software development practice in which automation tools are used to automatically and continuously build and test a software system as it is worked on, so as to ensure that the relevant outputs are always available and can be reproduced without effort.

this also produces an opportunity to evaluate the uncertainty that results from not knowing how different interventions overlap with each other. If the standardised interface is expressed including some notion of uncertainty or freedom, the possibility of monte-carlo type runs appears. To make this concrete, consider the following example: policy A provides a subsidy to many people, and the best estimate available is that 20% of households will adopt some intervention. Policy B is targeted toward areas with a high Index of Multiple Deprivation (IMD) score[8], and so it is easier to predict where the interventions will land, say on 50% of the target households. If the input to the integration layer described above preserves this uncertainty, it could construct a range of input scenarios for Transform from the cross product of these two distributions automatically, and produce a distribution on the whole system effect.

4.5 Making a map

Creating a framework for the automatic integration of multiple models may be too expensive to be justifiable; it would certainly imply a significant amount of re-implementation of existing work. Some of the same effects could be achieved using organisational processes rather than software systems; in particular, we would recommend making and maintaining a comprehensive central map of how different policy models are connected, and what factors they consider or are sensitive to. Such a reference would be a first step in a more systematic rationalisation, and would also be useful in its own right. This piece of work is a first step in that direction, but documentation is always easier to write whilst work is in progress than it is to construct after the fact, once people have moved around. As with our other recommendations, we suggest giving this responsibility to some people within DECC, and explicitly allocating time to the process of creating and maintaining the results.

5 What are the limitations of DECC's existing models for D3?

5.1 Sectors

In the models we have looked at, the domestic sector has the most detailed representation. This could equally be an artefact of the choice of models we started with, the current political focus, a reflection of where the most modelling effort is required, or even that it is easier to model because of the quantity and quality of data available / relative homogeneity of the sector. It seems relevant for D3 because the industrial and commercial sectors use a significant amount of energy in various forms, and are likely to contain some good candidates for all of the D3 interventions. In particular, there may be a lot of potential for demand shifting in the form of larger individual loads which can be controlled, which would need less control apparatus than a similar size of smaller loads from the domestic sector.

5.2 Technologies

These are some D3 technologies for which DECC's modelling abilities are currently limited in ways that may matter.

5.2.1 Heat networks

Heat networks seem like the biggest missing piece; as discussed in (2.2.2) their highly spatial nature makes them the hardest part of D3 to analyse at scale. Together with the fact that their value is tightly connected to other D3 measures - insulation through heat demand, heating technologies by nature, and electrical demand response, distribution and generation because of the co-generation aspect - a good analysis of, say, the national potential for heat networks, would have to tie together an understanding of all of D3 in concert.

---

[8] This example immediately shows how the cross-cutting aspects of D3 make analysis difficult - can substations in Transform be identified in a statistically representative way with respect to the IMD score of the Lower Super Output Area (LSOA) containing them?

5.2.2 Insulation

As far as we have been able to tell, none of DECC's models express the super-additive benefit of different kinds of insulation taken together (see (2.2)). The widespread use of SAP or SAP-like algorithms for the domestic sector precludes consideration of this effect, which has even been written into policies which relate subsidy to performance under SAP. The significance of this effect may of course be erased by other error terms, in particular the difficulty of determining accurate u-values (again, SAP's prominence may have some undesirable effects here; SAP u-value tables confuse the performance of new insulation with existing insulation), heating regimes and so forth.

5.2.3 Heat pumps

Heat pumps are more difficult to model than other single-building heating systems for all the reasons discussed above, and considered further in (2.2.1). At the moment, the only model which recognises these factors is ESME, which can consider both seasonal heat storage and the seasonal effect on performance to some degree. However, even ESME lacks a detailed picture of the types of buildings in which heat pumps are to be used (with respect to their thermal properties), and does not support an efficiency which varies with load anyway.

5.3 Temporal resolution

Correct pricing of electrical demand and network reinforcement costs ideally requires a high time resolution. Some of DECC's models (Transform, ESME, the DDM) can operate with hourly or sub-hourly time resolution, but the policy models we have seen do not produce forecasts with this precision. This may be because the evidence required to back up detailed forecasts of side effects is not available, but it does imply some possible systematic errors. For example, if the products policy models are generating most of their savings from improved refrigerators, their effect on peak load will be quite small. On the other hand, savings from lighting, or fuel switches in cooking may be weighted much more toward the peak.

Some understanding of the profile effects of different interventions must exist within (for example) Transform; perhaps this information and additional research or assumptions could be used to generate modifications to the typical load profile from the less detailed outputs that other models already produce; for example, a table breaking down changes in demand by end use could be augmented with typical time-of-use information to produce load-profile forecasts for use in evaluating system costs. This is discussed in more detail in (4.2) and (4.3).

5.4 Modelling causal relationships

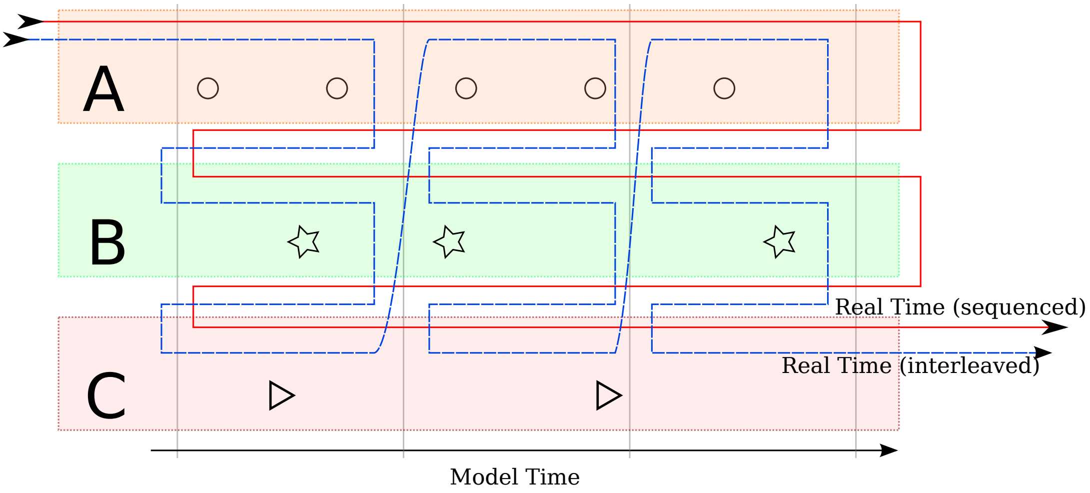

The image above shows two different ways of modelling time. Policies A, B and C are modelling interventions affecting circles, stars and triangles respectively. Model time travels from left to right, so in the modelled output the different interventions are interleaved in time. Consequently, in principle any side effects of each shape could affect any other shape drawn to the right.

The red line illustrates how if the models are run in sequence A B C, it is impossible for any star-shaped intervention to affect any circular intervention, or for any triangular intervention to affect either a circle or a star, in spite of the fact that in model time such effects are legal.

The blue line shows an alternative running order, in which the first year of A is processed first, then its consequences are used to compute the first year of B, and then the joint effects to compute the first year of C, before the second year of A and so on. In this interleaved sequence, the allowed causal structure is closer to the ideal; circles are still isolated from stars and triangles within each year, but inter-year exchanges are clearly possible. A time slicing that was finer still could allow the modelled interventions to represent all causally feasible interaction effects.

At the moment, DECC has a stack of models, some of which are partially interleaved, and some of which are entirely "downstream" from others. Although a careful case-by-case consideration may permit this kind of structure - say in the example above, an analysis may have revealed that circles are never affected by stars or triangles - an increasing number of models, and interventions being modelled, makes it more difficult to be certain about the correctness of this kind of assessment. The mutually-affecting possibilities for demand side intervention create a lot of potential for causal confusion in this way.

The smart meters model (6.6) is a case in point for this; it is run manually, after the EDM and all other policy models, and so cannot possibly be affecting any of their behaviour, or the price-demand relationship iterated between the EDM and DDM. Furthermore, no effect on load profiles (for example from Time of Use (TOU) tariffs) is produced for consumption by the DDM, which may be sensitive to such changes.

This difficulty can addressed by the use of smaller numbers of larger and more integrated models, like ESME and the NHM, which maintain a coherent internal state that can be operated on from different directions. In the NHM, for example, an arbitrary number of policies can be defined to take different interventions at different times, with guaranteed causal consistency (the effects of all preceding changes are always visible at any point). However, DECC's need to model a lot of different things in quite different ways, and to answer what-if questions at quite short notice is naturally in tension with the methodical, monolithic approach of using a few large systems like this, and there is always likely to be a need for smaller ad-hoc models. It may be that some benefits of both approaches could be kept by creating a way for multiple small policy models written using arbitrary tools to interact via a central "broker" that ensured a causally consistent shared state. We discuss this further in the next section.

5.5 Software engineering principles

There are some general issues about the way DECC's software is organised and connected together, and the problems which this approach produces for holistic understanding and consequently for modelling changes which have behaviours that must be understood holistically. Whilst it may be possible to use the existing set of tools to model some D3 scenarios, the range of Rumsfeldian "unknown unknowns" would be worryingly large, and a lot of manual work would be needed to integrate the results of all the different models together. There are some changes to the way DECC constructs and connects its models which we believe could ameliorate these problems, although they are not without their own attendant difficulties.

Because DECC's models are expected to persist for some time, be used by people other than their original authors, and to inter-operate to some extent, the modelling process is in part an exercise in software engineering. However, because few DECC employees come from a software background, and there is a lot of staff turnover between projects (and even departments), the tool set available for development is quite restricted. This causes quite a lot of excess complexity; normally in software development some effort is put towards paying off so-called "technical debt", the term used to describe design decisions which make things easier in the short term, but will increase the overall amount of work that has to be done when accommodating new requirements in the future.

Some technical debt is unavoidable, and there is a trade-off between the expected longevity and reuse of a piece of work, the importance of comprehensibility in the fine detail, and the cost of spending a lot of time on producing a robust design. The interconnected nature of D3 problems may demand a different trade-off between ease of model development and robustness of design than problems which are more separable, and so amenable to representation by a number of spreadsheets interconnected by small volumes of information. DECC's modelling system *may* currently be able to represent and analyse D3 measures effectively, but the risk that there is currently nobody in DECC with a holistic understanding of the models and how they may safely be connected always admits the possibility that some important details are being missed without anyone being aware of it. To some extent the purpose of this document is to produce such an overview, but ultimately we feel that the view of any external contractor will always be too limited to do the job as well as DECC could for itself, given the time and resources, and does not in the long run develop and maintain the knowledge where it is actually needed - within DECC.

Returning to the idea of DECC's style of modelling as an exercise in software engineering, we would recommend the consideration of some ideas from the world of software development and research.

5.5.1 Automation and integration tools

We would recommend that DECC look into more widespread use of automation and integration tools. In particular, DECC analysts are quite restricted in their range of tools and facilities for running software they have created. Making a lightweight, general-purpose scripting language available (Python, for example), and using it to ensure that the process of running a model is fully automated would make it easier to combine different models together. Even if different models are not interconnected in the best possible way, the ability to automatically run them together would shorten the exploratory feedback loop and make investigating system-level effects easier. Capturing the whole process for running a model in software also allows the use of continuous integration tools and remote servers to automate repeated runs, promotes reproducibility and transparency, and makes explicit the connections between models (for example, by eliminating the use of copy-paste or any manual approach to set up input data), reducing the risk of error and allowing better quality assurance.

5.5.2 Reproducible research

One current theme in academic research is the notion of (computationally) *reproducible research*, the idea that scholarly articles should be distributed along with a computer program capable of directly reproducing all the derived results in the work, and any raw input data. For a fine example of this, which also serves to illustrate how sequential programming avoids a lot of the issues around using spreadsheets for even relatively simple models, see [1], which reproduces the analysis from Reinhard and Rogoff's infamously erroneous work on interest rates and government debt [8], but without the errors.

5.5.3 Standardised interfaces

We have already described a specific example of how defined interfaces could be useful in (4.4); a more general adoption of this principle could help solve the problem of representing causality described in (5.4), although it would also imply a wholesale redevelopment of many existing models. If each model interacted with a shared description of the state of the world through a set of standard interfaces, interleaving them in time would become much easier and would not involve connecting the models to one another directly. Defining and using interfaces also promotes the explicit consideration of which interactions may be important, and makes disentangling the dependencies between parts much easier (because they are structured and controlled).

The EDM is an unusually good example of how this can be done using spreadsheets, in spite of the difficulties they present for modular design; its controlled use of formulae and automated ingestion of external data would make good reference cases for how to impose this discipline when designing a spreadsheet model.

Even where the tool chain or the nature of the work prevents the creation of formal (i.e. machine readable) interfaces to other models, we would recommend documenting interfaces as a separate task in all cases; for example, much of the work involved in producing this document could be avoided if DECC maintained a central description of how all its models are connected.

5.5.4 Comprehensive use of version control

Use of version control systems is another standard practice in the world of software, which is enabled by the use of sequential, text-based programming languages. World-class tools like git [9] provide an efficient, powerful way for multiple people to work on sets of text files, and maintain a complete history of changes and the reasoning behind them. Tools like this support experimental and offline working (as there is no risk of damage to the "authoritative copy"), and taken together with full automation of model execution would allow an easy way to see and store changes in model output over time, and to bisect history to find the cause of unexpected changes.

Dealing with the challenges of developing large software systems is not easy, but there are definitely some painstakingly discovered best practices to employ (and pitfalls to avoid) which DECC could benefit from; again, we believe this would best be done by maintaining the relevant expertise within DECC, rather than by contract. Although we have written about these ideas here because of the particular problems around modelling D3 - it is complex, and likely requires a complex model - the benefits would apply much more generally. They do, however, imply a change in approach; spreadsheets in particular provide a difficult environment for designing complex software, and impose a lot of cognitive overhead for systems above a certain size.

6 Model descriptions

6.1 Energy Demand Model

The EDM is an econometric model of the whole energy system, which forecasts large-scale energy demand as a function of some aggregate properties of the economy that are computed off-model. There is a programmatic link between the EDM and the DDM which is used to find a supply/demand equilibrium between electricity demand on the grid and the plant needed for generation.

The model is composed of a large number of simple econometric relationships (around 3500 equations, although some of these are simple aggregations or unit conversions), which break down Gross Domestic Product (GDP) across industrial sectors, and then relate per-sector GDP, fuel prices, population and so on to fuel requirements in each sector. These internal sector models are intended to predict business-as-usual fuel consumption in the absence of policies applied after a certain date; finally the effects of policies are then accounted for as changes in fuel consumption added to the forecast made in their absence. To allow for this, when the econometric models are retrained on new historical data the “backcast” effects of policies must be taken off the training data.

Figure 6 shows a ten-thousand foot overview of the EDM, and figure 5 shows some of the structure in the domestic sector in more detail; these are both given as illustrative examples.

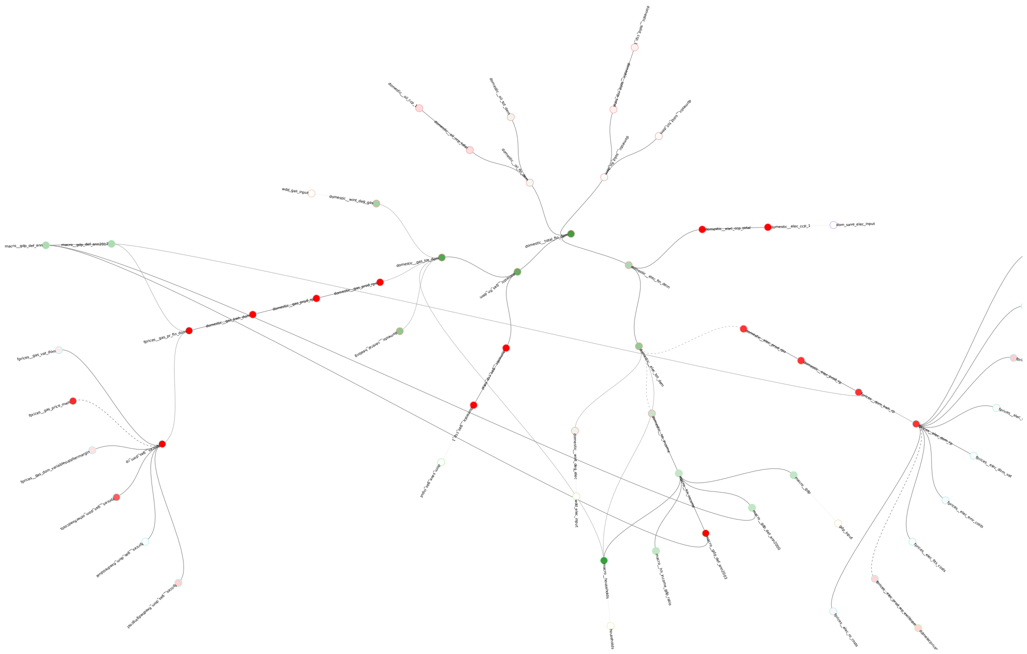

Figure 5: A schematic of domestic sector energy demand in the EDM. Each node represents an equation, with the central node being the demand forecast; edges radiate outwards from equations to their inputs. Nodes are also coloured by the sensitivity of the central variable on their value, red indicating anticorrelation and green correlation. A live version of the model which produced this figure is available as an appendix.

The fine detail of what systems are involved in the actual consumption of fuel to what ends is not explicitly represented in the EDM, but instead is implicit in the training data used to fit the econometric equations[9]. Because of this, the EDM cannot be used to predict the consequence of changes to the particular technology mix in the UK, although it can produce relations between aggregate figures like GDP and per-sector emissions. To find DECC’s representation of demand side detail in the EDM, we need to look to the off-model inputs.

---

[9] For a business-as-usual (BAU) forecast this is not a significant limitation, and some research [6] suggests it may suffice for any model presupposing global GDP growth.

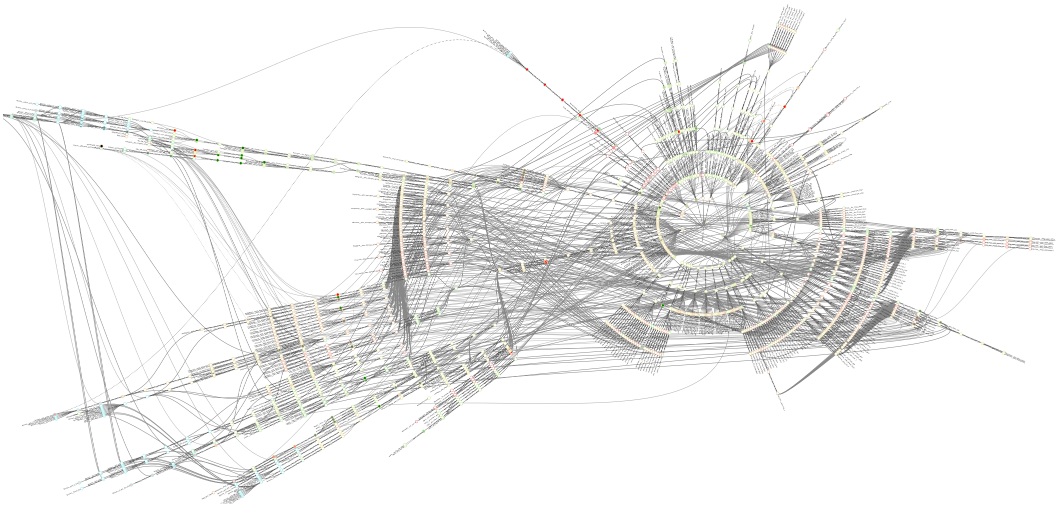

Figure 6: A schematic of total CO_2 emissions in the EDM. See figure 5 for an explanation; this figure is included to give a sense of the size of the model. Again, to get a useful understanding it would be better to look at the interactive version.

6.1.1 Model inputs

The EDM is the most well-connected model in the ensemble, which is appropriate considering what it is modelling, taking inputs from several other models. The largest number of connections are through the fuel savings forecasts.

Policy savings Policy savings are entered from a range of other models, as described in (3.1), and subtracted from the business-as-usual demands that the EDM would predict in their absence.

Electricity prices Retail electricity prices are taken from the DDM, and used to determine a supply demand equilibrium. The DDM interface updates `ddmelecprices`, and is then provided with `output_vars__ddm_elec_twh`.

The electricity demand model in the EDM is fairly price-inelastic (see figure 7), possibly because the training data does not contain any examples of what would happen if electricity prices rose or fell for an extended period. Furthermore any coupling of GDP to energy prices that may exist cannot be represented within the model. One small oddity here is that the base demand in the EDM does respond to price, albeit to a limited extent, but the demand offsets which are netted off the final consumption do not; conceptually this is a statement that all the energy savings made are in an inelastic part of overall demand.

Autogeneration The CHP model is used to forecast autogeneration figures for the EDM, which are mainly used to produce a better figure for grid transmission losses. Each sector has a variable indicating its CHP autogeneration; these values are summed to a total, which is divided into own-use and electricity exported to the grid. Own-use is then taken off the estimate of end-of-grid usage, which is in turn used to calculate the additional losses incurred in grid transmission. The CHP export share proportion is derived from DUKES, and is not forecast to change within the model. Own-use is *not* subtracted from demand here, because that calculation is performed within the DDM as avoided generation

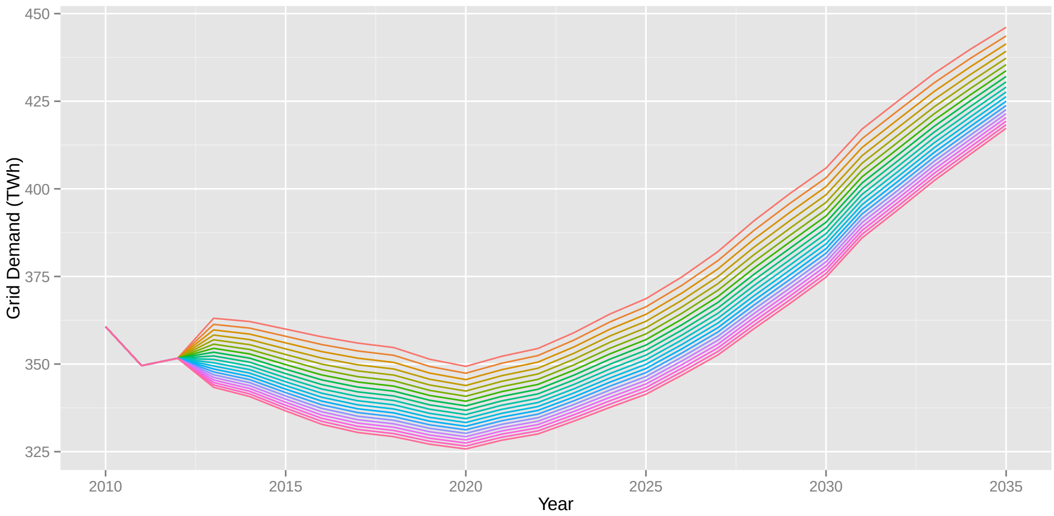

Figure 7: Electricity price-demand relationship in the EDM, where the price function has been adjusted from a tenth to double the base-case in the EDM run spreadsheet we used for this analysis - the lower estimate is with the doubled price. The same curve obtains at double the price, largely because the GDP forecast drives demand, and is itself independent of the price.

6.1.2 Model outputs

Electricity demand As already described, an electricity demand time-series is provided for the DDM to predict grid behaviour.

Heat demand The RHI projections use a forecast of heat demand produced from the EDM. At the moment we have not seen enough of this to know exactly what variables are involved.

Smart Meters The EDM is reported to have been used in the creation of a tailored baseline for the assessment of smart metering; this is presumably to reflect the fact that as demand changes the potential benefits of smart meters also change.

6.1.3 Policy outputs

Updated emissions projection One of the main outputs for the EDM is the Updated Emissions Projection document, published annually.

Interdepartmental Analysts Group Guidance The IAG guidance [4] is produced periodically and circulated to other government analysts. It contains the figures used to coordinate the value of of energy use and greenhouse gas emissions across government.

Prices and bills report The prices and bills report is another periodical output which summarises the absolute and distributional effects of energy policy on domestic energy bills.

Baselines for the CCC EDM outputs have also been used to produce baseline scenarios for the Climate Change Committee.

Price and GDP impacts for the Treasury Finally, the EDM is sometimes run at short notice to assess the effects of changes in GDP or fuel price forecasts on emissions and consumption for the Treasury.

6.2 Dynamic Dispatch Model

The DDM is a simulation of the electrical system from the point of view of the grid. An overview of the inputs and outputs of the DDM is given in figure 8; it illustrates that since most D3 behaviour occurs at the end of the distribution network, most D3 effects will appear through the inputs to the DDM.

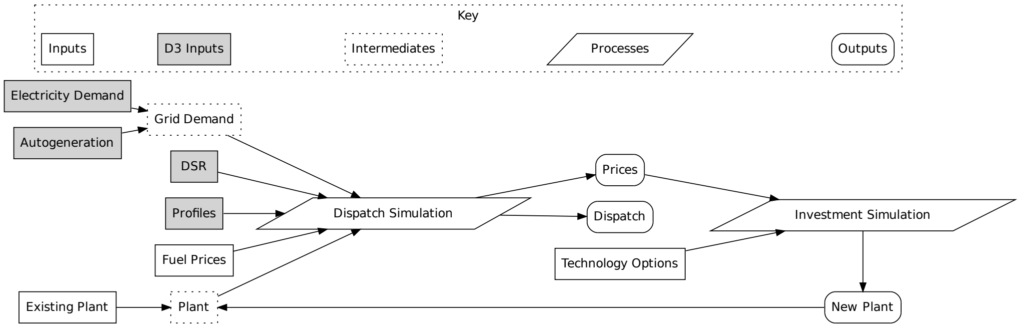

Figure 8: A simple overview of the DDM - normal rectangles are values computed off-model and rounded rectangles are outputs from the model. Coloured rectangles are inputs which could be varied when thinking about D3. This diagram does not correctly detail how the investment simulation and dispatch simulation are connected; in practice the dispatch simulation is used for two purposes, firstly for predicting actual prices and dispatch, and secondly to inform the simulated investor about her forecast for plant economics.

6.2.1 Model inputs

Many inputs are taken into consideration by the DDM; these are the ones most important in its consideration of the demand side.

Annual totals Annual aggregate demands are provided by the EDM to 2030, by MARKAL from 2040 and interpolated in between. The totals are reduced by 2.7% when not modelling Northern Ireland.

Load profiles Load profiles are taken from some sample days given by National Grid, broken down by some day types. These are taken as-is and do not change in response to anything on-model; they are simply used together with the annual totals (net of autogeneration) to produce load profiles which in turn produce the merit order and financials for new plant.

Autogeneration Autogeneration (demand met on the demand side, and excess power exported into the grid by distributed generation) is presented as a forecast timeseries, which is simply netted off the demand timeseries to produce the remaining demand which the grid must supply. Most of the autogeneration data is provided by the CHP model, which predicts future CHP capacity.

6.2.2 Model outputs

The only model systematically consuming values from the DDM is the EDM, as shown in figure 1.

6.2.3 Use with ESME

Because the DDM provides a good model of the costs and constraints for central generation it has also been used to produce inputs or augment outputs from ESME for specific pieces of modelling work. As an input to ESME, the DDM can be used to produce a projection for grid generation which ESME can be constrained to follow. Taking output from ESME, the construction of new plant can be fixed and the DDM simply used as a better tool for calculating electricity prices.

6.2.4 Policy outputs

The DDM is been influential on both policy analysis and some important documents; during our conversations we have come across the following:

- Electricity Market Reform Delivery Plans
- Gas Generation Strategy
- Hinkley Point C
- The UEP and IAG guidance

Apart from these specific instances, the model is frequently run for analysts in other parts of DECC to answer questions about the outlook for technologies they are interested in under different sets of assumptions, in particular:

- Office of Nuclear Deployment (OND), regarding nuclear
- Office of Carbon Capture and Storage (OCCS), regarding CCS
- Office of Renewable Energy Deployment (ORED), regarding offshore wind and marine energy

As we understand it, for these questions the model is not necessarily cycled with the EDM, but given the low price elasticity in that model it may not be important to do so.

6.2.5 Demand response

Changes to demand profile can be considered in the DDM by adjusting the load profiles which it consumes as inputs, but these cannot be adjusted on-model in response to any modelled outputs (prices, carbon factors and so on), so the effects of changes to peak on generation costs could be assessed by varying this input. This facility has been used for policy questions; firstly to assess smart meters, but the effects were quite small and it was unclear how much they represented modelling artefacts rather than likely consequences, and secondly as part of the electricity demand reduction pilot, to consider the effect of differential changes between domestic and non-domestic loads.

There is also a demand response option which can be selected in the merit order; unlike the DDM's model of normal plant, there is no facility for an investor to construct new demand response capacity (for example by fitting control systems to some existing industrial load, or by building energy storage systems). Instead, the amount of demand response possible is provided as a timeseries of price/power curves, so in any given simulated year the dispatch model may choose to pay for demand reduction according to such a curve. However, little use is made of this facility, and under the standard set of price and capacity assumptions (which do not forecast any future change) the model typically only uses demand response in extremis.

Some work has also been done modelling demand side response by inserting a special plant into the inventory which meets some demand with no fuel inputs, but apparently not in a setting where the model can choose to construct more DSR capacity. Indeed, it appears that the construction of new DSR is not exactly analogous to plant, as there is some upfront cost to making provision for DSR (like purchasing plant), but in the merit order there may be multiple prices at which that "plant" can provide some capacity.

6.2.6 CHP Module

A CHP module is under development to help the heat team in their analysis but it is not complete. It seems unlikely that it will be able to consider the spatial factors affecting the economics of heat networks.

6.2.7 Network costs

Although the DDM does not consider network costs during its simulation, they are presented in the cost/benefit analysis headline result that is typically produced. The network cost model used for this was developed by National Grid, and forecasts both transmission and distribution costs, and other costs associated with technical details around maintaining the grid's frequency.

In the UEP calculation (when iterating with the EDM), network costs are recovered by marking up the wholesale price output using a factor derived from historical data. Interestingly, the EDM also inflates grid demand by a transmission loss factor; there may be some double-counting here, but we are not certain.

6.2.8 Demand reduction

All demand reductions present to the DDM through the demand input; some demand-side generation is then subtracted as 'autogeneration'. Typically, the demand input is computed by the combination of the UEP's aggregate demand forecast and policy models' estimated savings relative to that baseline. The autogeneration input is computed by the CHP model, and represents a forecast for the amount of demand which is satisfied by on-site generation. This autogeneration forecast is only for certain industrial sectors.

6.3 Green Deal Household Model

The Green Deal Household Model is a forecast for uptake of Green Deal measures which takes some off-model market research and analysis of the housing stock, and then simulates expected changes to the housing stock under different adjustments of the policy's key "levers". Of all the active policies at DECC, after the RHI it has the greatest forecast savings entered into the EDM at `6.3` % of the total.

Figure 9 gives a very high-level overview of how the domestic sector spreadsheet model works; it forecasts uptake of a subset of possible Green Deal measures on the basis of some consumer research. The measures included are:

1. Internal solid wall insulation
2. External solid wall insulation
3. Cavity wall insulation
4. Loft insulation
5. Loft insulation with cavity wall insulation

The policy itself covers a much wider range of demand-side measures, including heating, lighting, glazing, and microgeneration[10], but these are not included in the spreadsheet model.

Although the GDHM spreadsheet only considers the measures listed above, the Green Deal IA [2] does contain forecasts associated with other measures like replacement boilers, energy efficient lighting and so forth. Boiler replacement in particular is considered to be entirely accounted for by Building Regulations, and those affected by Green Deal finance or any of the subsidies are presumed to be a subset of those which would have been replaced anyway. Exactly how the Building Regulations forecast feeds into the Green Deal analysis is unclear; the details of any subsidy ought to depend on the associated ECO/Affordable Warmth (AW) points

---

[10] All of these measures are eligible for Green Deal loan finance, but only some will attract subsidy under ECO, Carbon Emission Reduction Obligation (CERO) or Carbon Saving Community Obligation (CSCO).

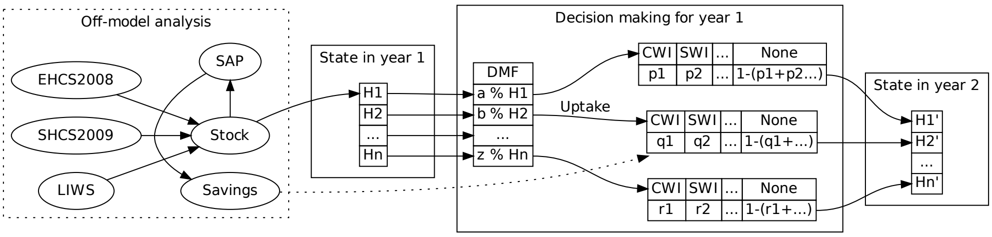

Figure 9: A schematic of one iteration of the GDHM spreadsheet model. Off-model, the English, Scottish and Welsh housing surveys have been processed into a representation of the housing stock; this is a distribution over around 3800 distinct "house types", which represent the starting state in the spreadsheet model. Along with this, a SAP-like calculation is performed to estimate the energy savings associated with each of the modelled measures. This done, the model then simulates each year by selecting a proportion of each bin of houses using a simple decision making frequency (DMF), and then a likelihood is for each alternative for each house based on the consumer utility function; this takes in a few bits of information, including the forecast bill savings. Every house that chooses to do a measure is then removed from consideration, and the remainder (those not deciding and those who have decided to do nothing) form the population for the next year.

score, which in turn will depend on the house's thermal performance and the rating of the boiler that is being replaced. For the two models to be correctly aligned, whatever criteria the Building Regulations model uses to determine in which houses the boilers are to be replaced should be replicated in the analogous Green Deal picture.

For those measures which are on-model, the IA includes a BAU projection of uptake, based on the GDHM's predictions with no subsidy or loans enabled. It is not clear whether the savings presented to the EDM are net of the BAU forecast, and if so where the BAU for insulation uptake enters into the EDM and how it relates to the base-case in the IA.

The IA also gives forecasts for the commercial sector, which are represented in the savings forecast to the EDM, but we have not seen the model behind these. Similarly, a BAU forecast is included which derives from a 2009 Element Energy report [5]. This forecast includes savings which are expected to be achieved from various technical improvements, some of which may be covered under Building Regulations (glazing and insulation) and products policy (lights, pumps and fans). Again, it is unclear how this base case is related to the forecasts from those models.

6.4 National Household Model

The NHM is not a specific policy model, but a framework for constructing a range of policy models. Consequently although it has several features which can represent some changes on demand side, it is only through using these features that it can actually help with any D3-type questions.

It contains an image of the UK's housing stock derived from the English Housing Survey, which model users can manipulate by selecting subsets of the stock and scheduling changes to be made in a discrete event-based simulation of time. The consequences of changes are predicted on a house-by-house basis using an implementation of SAP, which brings with it some benefits and limitations which we have already described.

Unlike most of DECC's models, the NHM's user controls are specified using a domain specific language, in which the user can describe scenarios by specifying rules for modifying the condition of selected houses. For example, listing 1 gives a very simple NHM scenario.

Listing 1: An NHM scenario which insulates 10,000 detached houses every year, and installs a new boiler in 4 percent of dwellings in London every 6 months. The random choices of subset may intersect; in this case the effects of previous interventions would be correctly accounted for.

```
(scenario start−date:01/01/2014 end−date:01/01/2040
   (policy
     (target
      name: wall−insulation
      group: (group.filter (is (house.built−form) Detached))
      exposure: (schedule.repeat interval:”1␣year” (sample.count 10000))
      action: (measure.wall−insulation type:Cavity resistance:0.2 thickness:50))

     (target
      name: new−boilers
      group: (group.filter (is (house.region) London))
      exposure: (schedule.repeat interval:”6␣months” (sample.proportion 4%))
      action: (measure.standard−boiler fuel:MainsGas efficiency:0.9)))
```

The scenario language is quite flexible, and allows the use of any of these demand-side measures:

- New heating systems, including:
  - Standard and condensing boilers.
  - Air or Ground source heat pumps (although with the limitations of SAP's standard model).
  - Electrical storage heaters.
- Various kinds of insulation:
  - Internal, external and cavity wall.
  - Floor and loft insulation.
- Behavioural changes, by modifying:
  - The heating temperature.
  - The heating schedule.
- Solar thermal hot water

All of these can be combined into packages, and selected between on the basis of arbitrary functions of the values the model computes (so minimum NPV, or a weighted random behavioural model, or maximum emissions reduction/cost). Causal consistency is maintained between interventions, so in a scenario like listing 1 if one intervention insulates some houses and another installs new boilers, they will each take the other's effects into account (so the benefit of installing a new boiler in an already insulated house will be reduced, for example). The model can also be used to run scenarios with random behaviour repeatedly, to discover the distributional effects that may be produced, although this is done using a brute-force monte-carlo method and so is computationally inelegant.

The model does not represent anything upstream of individual dwellings, and so provides no understanding of reinforcement costs or the relationship between price and demand for different fuels. Also, as it uses a SAP-based energy calculator the best temporal resolution available is monthly - there is no facility for the production of load profiles. These values could be computed off-model using other tools (Transform and the DDM, for example), but not in a way that would allow any feedbacks - the causality problem described in (5.4) prevents this. The model is not spatially representative (because the English Housing Survey on which it is based does not provide enough information), and so cannot usefully model the costs of technologies for which spatial effects are important like district heating.

At the moment the NHM does not provide any machine readable interfaces for integration into other workflows, but the concise form of the scenario language and strict reproducibility of its outputs means this would not be difficult to incorporate. However, because DECC does not make extensive use of automation at the moment, this is not currently very limiting.

6.5 CHP Model

As part of its CHP-QA role, AEA maintains detailed information on CHP plants subject to QA in the UK. Based on this information AEA has developed a cashflow model for CHP plants. This is run by AEA on behalf of DECC to explore the potential for CHP under varying scenarios expressed through the model's parameters. These cover the following factors:

- Interest rates
- Tax rates
- Plant efficiencies
- Plant capital costs
- Plant operating costs
- Fuel prices and emissions factors for gas and electricity
- Export prices for electricity
- Climate change levy rates

Because of the confidential nature of the CHP-QA programme information, AEA summarises the data into categories ("tranches") and provides DECC with modelling results at this level. Each tranche contains a specified number of sites in wide range of economic sectors, and these are further differentiated in terms of plant type, fuel, capacity, and whether the emitted CO_2 is traded on the EU Emissions Trading Scheme (ETS). Note that this model does not cover CHP for use in district heating networks - the domestic sector is not included, and heat produced from CHP is assumed to be used on site. DECC specifies a number of scenarios using the above parameters, and AEA runs its cashflow model internally, sending DECC a summary of the results. DECC then runs the model itself, this time on the tranche-level summary data, to perform monte-carlo type analyses against selected parameters.

The results from these modelling processes are used to explore the interaction between CHP plants and various policy frameworks such as Carbon Reduction Commitment, EU ETS, Capacity Mechanism, RHI, ROCs, CCL and Carbon Price Support. They also define the autogeneration capacity used in the DDM, which in turn affects the carbon emissions factor for electricity used in the EDM.

6.6 Smart Meters Model

The Smart Meters Model is an Excel spreadsheet based tool developed by Mott McDonald for DECC in 2008 and 2009 and subsequently further developed both in-house and by external consultants, including PA Consulting and Redpoint Energy/Baringa Partners. The cost-benefit analysis (CBA) model is maintained and owned by DECC. Its main purpose is to generate cost/benefit analyses for given roll-out profiles and associated assumptions to support the internal appraisal of policy options and the external publication of updates to the Smart Meter Impact Assessment. The model covers the roll out of both gas and electricity meters. It is frequently rerun to support analysis, hypothesis testing and options appraisal, mainly within the Smart Meters Implementation Programme.

Cost and benefit information is collected from a variety of external sources and stakeholders, and for the most part is converted to per-meter values. These are then aggregated based on the number of meters installed in each year of a rollout scenario. A range of adjustable assumptions are included in the model, changes to which influence the NPV of the Smart Meters Implementation Programme as calculated by the model.

Costs and benefits are collated from external sources and converted to per-installed-meter values. These are then aggregated based on the number of meters installed in each year of a rollout scenario, and this is used to examine the cost/benefit characteristics over time, of different scenarios.

6.6.1 Costs

Costs are split into the fixed costs for central functions required for the rollout (e.g. Data and Communications Company (DCC), comms hub, etc.), and the variable costs of installing different numbers of meters. Central costs are taken from the Ofgem RFI. A more detailed set of costs are presented in the impact assessment, but we have not had time to read it in detail.

6.6.2 Benefits

The main sources of benefit in the model are via avoided costs in generation, transmission, distribution, and customer service. These arise from the following effects, and access to half-hourly data is assumed to be available where required:

- Overall demand reduction (saved energy and carbon, which itself valued)
- Peak load shifting (reduced required capacities for generation, transmission and distribution)
- Better infrastructure investment planning (more optimal scheduling of asset management)
- Better infrastructure management (cheaper fault management)
- Reduced costs to serve per customer (automated billing, reduced complaints).

Overall demand reduction (saved energy and carbon) is assumed occur as a result of better energy management by end-users, facilitated by in-home displays(IHDs). The assumptions about demand reduction effects here are taken from a number of previous studies in the US and UK (these are referenced in the Impact Assessment). The model does not explicity represent price elasticity of demand; rather, an average saving is assumed and applied on a per meter basis.

Peak load shifting is assumed to result entirely from TOU charging for energy. This effect is represented in the model as a proportion of the peak load being shifted to off-peak time. The values used come from internal analysis at DECC. The benefit of peak load shifting is quantified as the sum of the values of avoided generation, transmission and distribution costs associated with the shift.

Assumptions regarding reduced infrastructure investment costs and cheaper fault management are static values, come from the original consultancy by Mott McDonald; the expected benefits to networks were found to be in line with an ENA report in 2012. We have not yet found out where reduction in costs to serve end-users comes from. Energy and carbon prices and carbon emissions factors are all taken from IAG (which in turn comes from the UEP).

6.6.3 Interactions with other models

The only interactions are with the EDM: the Smart Meters Model takes total gas and electricity demand from domestic and non-domestic users from the EDM. Based on these values it sends back TWh demand reductions via the EDM's policy savings interface. This interaction is handled manually. Any policy ordering questions are handled by the EDM team, and are not visible to the smart metering people.

6.6.4 Possible weaknesses with respect to D3

The model assumes that significant load shifting occurs as a result of the introduction of smart meters: this effect contributes to the benefit side of the CBA. However, these effects are not accounted for in the way the DDM simulates required generation capacity and investment patterns. Accounting for these effects might change the results of the DDM significantly, which could in turn feed back into the smart meters model via the EDM's total demand values by sector.

The model doesn't contain any representation of meters, networks or generators, nor is it connected to any models that do. It simply collates a set of assumptions about the costs and benefits of smart meters, translates them onto a per-meter basis, and allows multiplication by n meters over a multi-year timeseries. It is more like a compilation of external assumptions than a model.

6.7 ESME

ESME is a linear optimisation model of the whole energy system, which can be used to generate a cost-optimal path for the deployment of technologies to meet demands, insofar as the cost can be considered as a constrained linear function[11]. Figure 10 gives a high-level picture of how ESME works; it takes a set of predictions about the final demand for different energy services over time, and produces a best-cost plan to satisfy those demands subject to constraints on what can be built.

Like most mathematical programming models for such systems, ESME represents the energy system as a set of flows with constraints providing the conservation rules[9]. Each technology provides a certain capacity and efficiency for converting some form of energy (such as a fuel) into another (such as a demand), and the solver finds a cost-optimal balance of flows down each path to meet the required output. Figure 11 shows a simplistic representation of how domestic need for heat is represented in the model. Change in the technology mix over time is simply another set of constrained flows, which can be provided by the end user.

ESME has a lot of flexibility in terms of how it can be constrained and what it can represent. Essentially any conversion technology can be included so long as its behaviour and cost are linear (or piecewise linear), and any of the flows in the model should be amenable to arbitrary off-model constraint.

However, it also has a degree of rigidity which would be difficult to escape or avoid. In particular, the model has a very limited spatial understanding, dividing the country into a few regions between which some simple (again linear) connections can be used to represent energy flows. There is no representation of the electricity distribution network, and the generation and transmission systems are quite simple; for this reason the DDM has been used together with ESME, at least for Heat Strategy work, to produce a more reasonable plan for generation, based on ESME's forecast of electricity demand. Finally, it is focused on producing a cost-optimal plan subject to its constraints, rather than simulating a realistic range of different scenarios; as a

---

[11] Whether or not the choice of a linear model causes *significant* omissions in what ESME can consider is out of the scope of this document, but it is definitely possible; linear relaxations of nonlinear problems can have optimal solutions arbitrarily far from a nonlinear optimum, and quite a lot of things in the energy system have a degree of non-linearity - it is hard to purchase very small quantities of nuclear power, or to install a small amount of widely distributed district heating, for example.

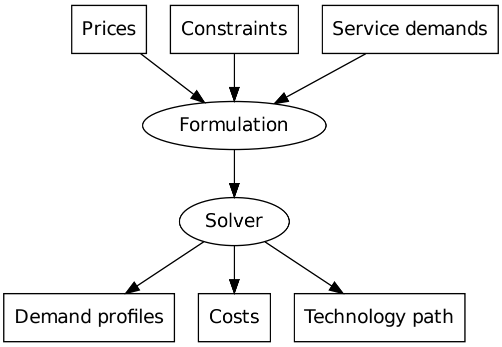

Figure 10: An overview of ESME. The user enters three rough categories of information: (a) demand for services (end uses of energy), (b) prices for material inputs and construction of new technology, and (c) constraints on flows of energy, build rates of technologies and so on. These are used to produce a formulation suitable for the solver, which produces a solution minimising a cost function. Finally the details of this minimal solution are produced, which include the suggested build order for new technologies, and the resultant energy and money flows.

result it may be a better tool for finding plans that meet high-level goals than for predicting the effects of policy in a world full of opposing interest groups and less than perfectly rational actors.

On the demand side, the model represents most of the end-users of energy in the UK, and can consider the uptake of a wide range of different conversion technologies.

6.7.1 Demand reduction

Any reduction in end use of energy must be exogenous to the model; however, ESME can represent a wide range of technologies which would produce demand reductions. Both domestic space heat and hot water are end uses, and there are a lot of different conversion technologies available to meet these demands, including

- Standard gas boilers
- Solar thermal
- Electrical heating, including heat pumps
- Heat networks

Because the model is not very spatial, those which are influenced by local factors will be less accurately presented; in particular, any treatment of heat networks seems hard to do at such a high level, especially when considering local electricity generation. Solar systems and heat pumps also have local effects (and even condensing boilers, to a really small degree), but these may be better represented by a few big regions.

Other demands have slightly simpler representations, but they are probably sufficient; lighting demand is essentially controlled by floor space and met by one of three types of lighting technology; other appliances are lumped together.

Demand reduction by insulation is less well modelled; it is only available for the domestic sector, and only as a rough classification of dwellings into a few different types; from a thermal standpoint there are three kinds of house, at least with the standard inputs. The model may insulate dwellings by paying some unit price to convert low thermal performance houses into medium or high thermal performance. This is clearly a large simplification, although we do not have the evidence to say whether it is fair or not here.

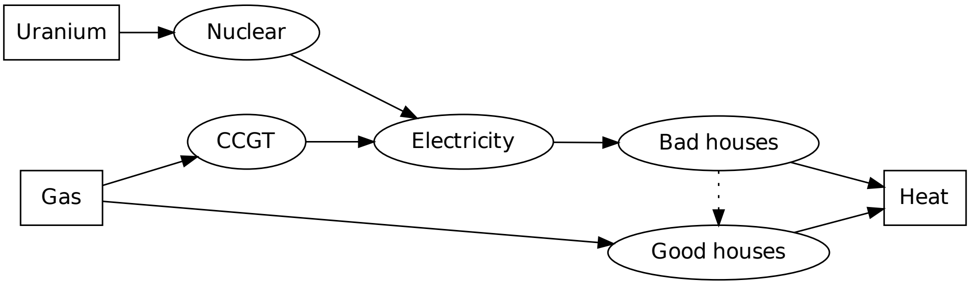

Figure 11: A simplified example of how ESME represents a subset of the systems it considers. Each oval node represents (effectively) an energy conversion technology, and each edge shows a path along which energy can flow. Limitations in capacity restrict the maximum flow on each edge. This shows how even things like the housing stock are considered a conversion technology which satisfy the total domestic demand for heat, given as an input. The dotted edge between good and bad houses indicates how change in the technology mix is represented as another kind of flow. A more correct diagram would show how time is represented as a series of connected variables; there really should be one copy of the entire diagram for each year, with causal relationships between years expressed as some additional constrained flows.

6.7.2 Demand response

Although the model cannot represent demand side response for its end use demands (as they are immutable inputs which must be satisfied), it does have one of the best treatments of sub-day demand from all of these models, and it does admit electrical storage as a technology for purchase. Storage and demand response are in some sense equivalent, and the only other model which appears able to consider sub-day resolutions is the DDM which has a very simple picture of demand. However, because of the model's lack of understanding of distribution costs and simplified view of generation costs, it is hard to say how its optimum cost relates to the actual cost of implementing a particular solution; it seems like ESME can produce load profiles whose system costs should be evaluated using other models.

6.7.3 Distributed energy

ESME's energy options include a range of technologies which *can* be distributed (solar PV, wind, CHP and heat networks, electric vehicles), and these are represented at sub-day resolutions allowing understanding of their relationship to peak load. However, as with the demand reduction technologies the spatial aggregation required by the model and the lack of any distribution network suggests a need for some off-model work to determine the significance of these factors.

Heat networks again seem like the biggest limitation here; the model can choose to build some total heat network capacity, but the financial viability of doing so really depends on the specifics of local heat and power sources, sinks, and distribution across the country. These would all have to be reduced to a few regional piecewise linear functions, and in any case imply the existence of a detailed local model which we have not seen.

6.7.4 Peak effects and network costs

As already mentioned, the model does have sub-day temporal resolution and a notion of the winter peak; however, its generation cost model is definitely a simplification (which is why the DDM's is preferred), and its distribution network model is non-existent. For either of these, external models (like the DDM or Transform) could be used either to produce sensible constraints or costs for input, or to replace that aspect of the cost function when pricing output; naturally, replacing the cost function post-hoc eliminates the solver's optimality guarantee.

6.8 MARKAL

We have not examined MARKAL; however, we understand it to be a mathematical programming model (a linear program, or similar) akin to ESME or the forthcoming UK TIMES. It seems to have limited use within DECC at the moment, and may be on the way to being phased out entirely.

6.9 Transform

Transform is a new model of the electricity distribution network which is not currently used for impact assessments by DECC, but is under evaluation as a way of solving some of the problems around reinforcement costs. At the moment it has been used to consider reinforcement costs for the following policy groups:

- OLEV, with respect to the possible deployment of electric vehicles
- ORED, to understand the impacts of small-scale PV and wind
- The heat strategy group, to consider the effects of electrified heating with heat pumps

The model is based around a description of the population of substations in the UK; about 100 archetypal substations are used in a weighted combination to represent the end of the distribution network. Each archetype describes a substation in terms of its electrical properties (maximum capacity, thermal rating and so on), the houses which are attached to them (their demand profiles for different types of day, and some other characteristics), and their connection to higher levels of the distribution network.

Given these archetypes, the model can forecast the load at each simulated substation, and on substations above them in the distribution network, and determine which substations are overcapacity and will require reinforcement. Reinforcement costs are then determined by optimising over a choice of interventions using one of two investment strategies (a shorter or longer term approach), and an output cost/benefit is produced from a network operator perspective and a wider social perspective.

To model possible future scenarios, the user can enter a time-series of technologies which change demand at each substation, for example by adding electric cars and solar PV to the system; under differing scenarios, different reinforcements will be proposed to ensure no substations are overcapacity. Figure 12 shows a cartoon of how this process works.

Interestingly, to allow a proper assessment of the social cost-benefit of a scenario, Transform includes a grid dispatch model with an inventory of plant, similar to that in the DDM. However, it does not match the DDM's ability to simulate investment in new plant and retirement of old plant in response to demand.

Transform's list of interventions includes electric vehicles, heat pumps, end of network PV, and PV, wind and biomass generation connected at the intermediate levels of the distribution network. It can deploy a wide range of solutions to reinforce the network, from predictable capacity increases to some more novel options, including demand response and energy storage.

CHP is one D3 technology which Transform does not currently represent, in particular for district heating. There is apparently work in progress to add a CHP representation to the model, but its scope is not clear.

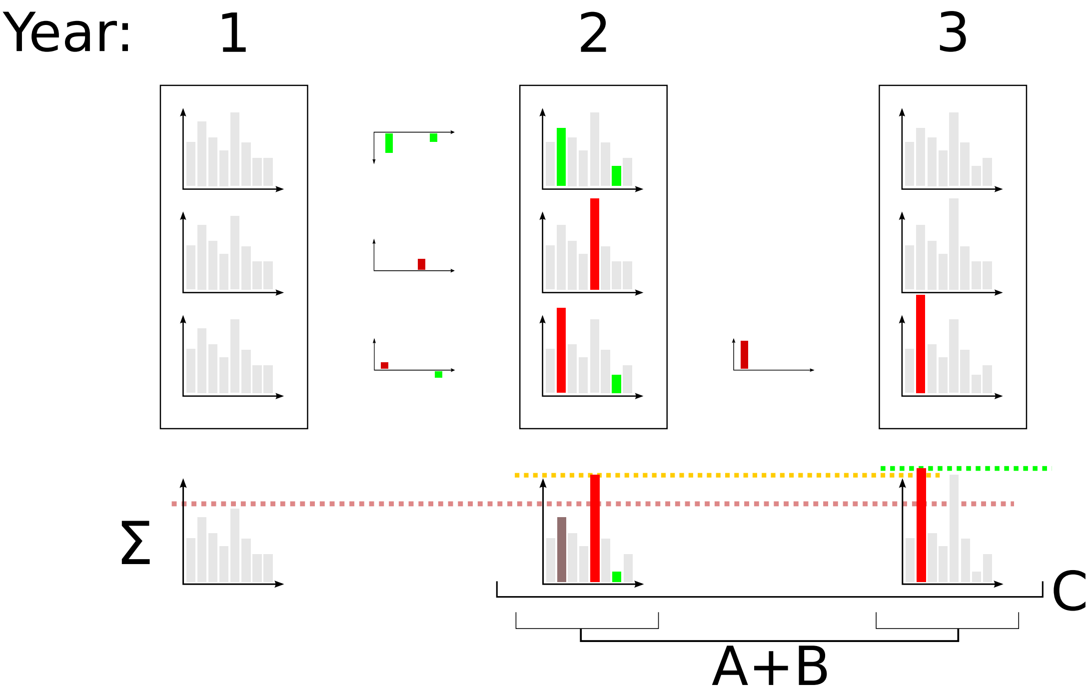

Figure 12: A schematic of a Transform scenario. The rectangles here show the state of a substation in terms of some load profiles for each of the three houses connected to the substation. Below each rectangle is a figure showing the summed load on that substation - this is the detail which matters as far as the substation's capacity rating is concerned, drawn here as the dotted red line. In year 1 of the scenario, some intervention is made which changes the load profiles of the connected houses, which is drawn in between years 1 and 2 as small deltas. The intervention causes the substation to breach its peak capacity rating, so reinforcement is required. The second detail shown is how the choice of period over which to optimise reinforcements is significant; if the substation were repaired for the minimal cost in year 2, using reinforcement A which raises the capacity to the yellow line, then the second intervention in year 2 would breach the capacity threshold again requiring further expenditure on reinforcement B to reach the green line. Using a two-year window instead allows the selection of reinforcement C whose cost exceeds A but is less than A+B.

6.10 Products Policy

At `11.5` %, products policy is one of the most significant current impacts on demand side forecasts. The following table shows the fuel/sector breakdown for the policy's savings; across sectors, it mostly reduces electricity consumption, and (as a result of reduced waste heat from electrical appliances) increases gas consumption, with the greatest effect on the domestic sector. Products policy targets devices in the home; from other DECC publications [3], most of this is due to improvements in efficiency of white goods, lighting and other electrical goods like televisions.

Fuel	Residential (TWh)	Commercial (TWh)	Industrial (TWh)
Electricity	470.11	198.43	40.12
GAS	-224.65	-11.74	-2.1
OIL	6.71	10.04	0.31

Products policy is mostly concerned with changes in electrical appliance efficiency over time; rather than a single model, the impacts of products policy are produced by a large number of small models, one for each kind of product. We have only seen one of these models and unfortunately not had contact with its author.

The model is quite a large spreadsheet, and without much documentation we have not been able to form a complete understanding; however, it seems that most of the complexity is related to modelling the replacement of the installed base of the technology in question over time, and that the costs and benefits associated with interactions with the energy network are produced in a simple way with reference to the IAG guidance document (which is in turn produced using the UEP).

6.11 Building regulations

Building regulations set the energy performance standards that have to met when building work is carried out. This includes the construction of new buildings, the extension of existing buildings, conversion work arising from change of use, renovation of the external building fabric, and the replacement of controlled fittings (like windows) and services (such as boilers or Heating, Ventilation and Air-Conditioning (HVAC) systems).

Each time the Building Regulations are updated, DCLG produces an IA which predicts the change from that the revisions will cause, relative to the prior regulatory environment. The cumulative effect of each successive tightening of regulation is what DECC uses for a delta to the EDM. Energy modelling is carried out using the most recent version of SAP for homes and of the Simplified Building Energy Model (SBEM) for non-residential buildings, under the usual assumptions of standard occupancy[12], fuel costs, and carbon intensity. For some representative buildings of various sorts, these methods are used to determine the building performance under (a) the minimum standard allowed prior to new regulation and (b) the same following the new regulation. These figures are scaled up using national build and replacement rates and the lifetimes of different technologies, and combined with Green Book and IAG supplementary guidance to produce the final set of costs and benefits.

These impact assessments have become more refined with each iteration with more robust sensitivities for build rate, build mix, compliance levels etc. (including for new homes looking at discrepancy between design and as built performance), better counterfactuals, better distributional, small firms and rural impacts and so on.

6.11.1 Suitability for D3 modelling

Whilst the modelling approach taken here suits the original purpose well, there may be some D3-related omissions around the way it is integrated with other work. Building Regulations are clearly a D3 concern - because they lower bound the thermal performance and heating efficiency in new builds and in work done on existing buildings, they reduce demand for heat and power. This should affect other D3 policies in that:

1. The scope for further demand reduction is reduced - as Building Regulations slowly improve all the boilers in the country when they break down, the scope for savings from more efficient heating is reduced.
2. The savings available from competing demand reduction technologies are reduced - insulating a house with an efficient boiler is less valuable than insulating one with an inefficient boiler.
3. The amount of demand available to shift may be reduced, depending on whether the associated reductions are in demands that are suitable for DSR.
4. The distribution of heat demand in the country will slowly change, as thermal efficiency changes; this may change the economic viability of certain heat networks.

The size of the impact is very large - Building Regulations own the bulk of the reductions passed into the EDM (`39.8` %). The following table shows the breakdown of energy savings attributed to the Building Regulations and used as inputs to the EDM: the overwhelming majority are in domestic gas usage, due largely to gas boiler replacement and increased efficiency in new buildings.

---

[12] The use of standard occupancy is slightly erroneous for DECC's purposes (although reasonable for the regulation), because DECC are ideally looking for an accurate national figure; however, a comparison we have done using the NHM gives a difference of about 1.5 TWh/year for domestic consumption between standard occupancy and EHS surveyed occupancy, which is less than 1% of the total.

Fuel	Residential (TWh)	Commercial (TWh)	Industrial (TWh)
Electricity	44.86	68.18	74.71
GAS	1081.88	128.09	158.68
OIL	46.48	18.33	40.42
Renewables/bio fuel	3.66	0	0
Solid	13.63	0.21	3.85
Renewables	0	0.64	2.07

Firstly there are two quite specific issues which may be significant for DECC's use of the model with respect to D3;

- Building Regulations are slightly different in the devolved regions, and their effects have to be accounted for by scaling (although this may not be significant).
- The impact assessments are based on the conservative assumption that work will be done to the minimum standard[13]; in fact, sometimes work will exceed the minimum standard for one reason or another. Whilst this makes sense in a cost-benefit assessment of the regulation, the resultant undercounting may be distorting the policies' effects for DECC's purposes.

Aside from these, as elsewhere, most of the difficulties that may exist here arise from sequencing the impact of policies, and from determining their effects in combination.

The effect of appliance efficiency is convoluted; it seems that the products policy models own the additional cost to space heating that they produce, but in conversation with DCLG they indicated that changing thermal gains are taken into account in their forecast; we expect that this is simply the improvement in products up to the date of the regulation, but if it does incorporate future changes then there is a double-counting here. Whether or not this error exists, there does not appear to be any integration of the effects of Building Regulations within the products policy modelling, so the same double counting looks to be occurring on that side (the energy for heat displaced should slowly fall as buildings are improved in accordance with regulation).

Apart from appliance efficiency, the effect on the housing stock must be taken into account. This is difficult because there is no standardised approach; it seems that, for boilers, the presumption is usually that all boiler replacements accrue to the building regulations model (the Green Deal IA takes this line, as does some Affordable Warmth modelling we have seen). This seems reasonable, so long as (a) no policy causes the replacement of a boiler which would not otherwise have been replaced, (b) that the efficiency and lifetime of the replacement are no better than the Building Regulations presume, and (c) that no policy causes a failed gas boiler to be replaced with some other device (for example a biomass boiler).

The other major impact is on insulation; the Green Deal is an obviously relevant policy here, whose insulation related effects seem to be measured against a baseline scenario with no intervention. In this case, there is no double-counting of the benefit of renovations, so long as (a) the baseline forecast by the Green Deal behavioural model agrees with that used for Building Regulations work, and (b) the baseline insulation that is simulated is the same as that predicted by the Building Regulations work (i.e. it meets but does not exceed the minimum standard).

Since the Building Regulations forecasts are based on large-scale trends with a lot of inertia, having the model causally upstream of several other models seems reasonable, even across the whole span of time being considered. However, the complexity associated with avoiding double counting its effects serves as a fine illustration of how a centralised representation of policy effects like that suggested in 4.4 (or at the very least a comprehensive map of which effects belong where) would make thinking about D3 easier within DECC.

---

[13] Personal communication with Paul Decort at DCLG.

6.12 Feed-in Tariffs

As mentioned in the introductory section, we have not had any time with the FITS model or its owners, so this description is necessarily limited. The model is implemented in Excel, and is at (or beyond) the maximum size and complexity for which that is sensible. It is mainly concerned with simulating the economic behaviour of several types of investor, who may procure some of a range of generating technologies in response to feed-in subsidies and the allocation of Renewables Obligation Certificates. Naturally, such an approach is very sensitive to the underlying economic assumptions about the rationality of its investors and the macroeconomic trends; for example, a major recession might dramatically increase the attractiveness of Feed-in Tariffs, as the rates attainable on other investments fall.

The model represents a range of technologies at different scales:

- CHP, from domestic micro CHP to 25kW gas engines or large biomass
- Wind, from 1kW to 5MW
- Hydroelectric power, from <15kW to 5MW
- PV, from domestic 2kW to 5MW
- Energy from waste plants, up to 0.5MW

As far as we are aware, the model does not take into account network reinforcement costs associated with connecting generation to different tiers of the distribution network - it certainly does not have a sophisticated understanding like Transform would provide. Its effects are however felt by the DDM, which takes the forecast useful generation from the FITS model as autogeneration, which the grid can avoid producing. There is no synergistic interaction possible between the FITS model and storage or DSR technologies which might allow better alignment of supply with demand, as none of the models responsible for such things are connected to it.

The FITS model does predict uptake of CHP, which is provided to the DDM, but apparently not to the EDM, unlike the CHP forecast produced by the eponymous CHP model. As far as we can tell there is no connection between the FITS view of CHP and the CHP model's view; it seems possible that there ought to be, as they are both models of investor behaviour, and it seems unlikely that they are considering entirely independent sets of investors.

7 References

[1] Vincent Arel-Bundock. Replication of Reinhart-Rogoff “Growth in a Time of Debt”, 2013. `http://nbviewer.ipython.org/github/vincentarelbundock/Reinhart-Rogoff/blob/master/reinhart-rogoff.ipynb`.

[2] DECC. Final stage impact assessment for the green deal and energy company obligation. Technical report, DECC, 2012. `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/42984/5533-final-stage-impact-assessment-for-the-green-deal-a.pdf`.

[3] DECC. Estimated impacts of energy and climate change policies on energy prices and bills. Technical report, DECC, 2013. `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/172923/130326_-_Price_and_Bill_Impacts_Report_Final.pdf`.

[4] DECC. Valuation of energy use and greenhouse gas (GHG) emissions - Supplementary evidence to the HM Treasury Green Book on Appraisal and Evaluation in Central Government. Technical report, DECC, 2013. `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/254083/2013_main_appraisal_guidance.pdf`.

[5] Element Energy. Uptake of energy efficiency in buildings. Technical report, Element Energy, 2009.

[6] Tim J. Garrett. No way out? The double-bind in seeking global prosperity alongside mitigated climate change. *Earth System Dynamics*, 3(1):1–17, 2012. `http://www.earth-syst-dynam.net/3/1/2012/`.

[7] Kiwa GASTEC. SEDBUK boiler efficiency database. Technical report, CRE, 2010.

[8] Carmen M. Reinhart and Kenneth S. Rogoff. Growth in a Time of Debt. Technical report, National Bureau of Economic Research, 2010. `http://www.nber.org/papers/w15639`.

[9] Linus Torvalds. The git version control system. `http://git-scm.com/`.

[10] Howard Wainer. The most dangerous equation. *American Scientist*, 95(3), 2007.