

A decision support system for addressing food security in the United Kingdom

Martine J. Barons¹  | Thais C. O. Fonseca^{1,2}  | Andy Davis^{1,3} | Jim Q. Smith¹

¹Department of Statistics, The University of Warwick, Coventry, West Midlands, UK

²Federal University of Rio de Janeiro (UFRJ), Rio de Janeiro, Brazil

³Coventry & Warwickshire Local Enterprise Partnership (CWLEP), Coventry, UK

Correspondence

Martine J. Barons, Department of Statistics, The University of Warwick, CV4 7AL Coventry, UK.
Email: Martine.Barons@warwick.ac.uk

Funding information

EPSRC, Grant/Award Number: EP/K039628/1

Abstract

This paper presents an integrating decision support system (IDSS) for food security in the United Kingdom. In ever-larger dynamic systems, such as the food system, it is increasingly difficult for decision makers (DMs) to effectively account for all the variables within the system that may influence the outcomes of interest under enactments of various candidate policies. Each of the influencing variables is likely, themselves, to be dynamic subsystems with expert domains supported by sophisticated probabilistic models. Recent increases in food poverty in the United Kingdom have raised the questions about the main drivers of food insecurity, how this may be changing over time and how evidence can be used in evaluating policy for decision support. In this context, an IDSS is proposed for household food security to allow DMs to compare several candidate policies which may affect the outcome of food insecurity at the household level.

KEYWORDS

Bayesian multi-agent models, causality, coherence, decision support, graphical models, integrating decision support systems, likelihood separation

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. Journal of the Royal Statistical Society: Series A (Statistics in Society) published by John Wiley & Sons Ltd on behalf of Royal Statistical Society

1 | INTRODUCTION

This paper gives a proof of concept practical application of the recently developed statistical integrating decision support system (IDSS) paradigm. An IDSS is developed for policymakers concerned with deciding between candidate policies designed to ameliorate household food insecurity within the UK context of rising food charity use. This paper starts with a brief overview of food security in the United Kingdom and then summarises the IDSS statistical framework. Next, the paper describes how the framework was used in developing decision support, including the role of policymakers, assumptions and sources of data. Finally, several scenarios and policy decisions are explored and the efficacy of the approach discussed.

1.1 | Food security

Food security exists when all people, at all times, have physical and economic access to sufficient, safe and nutritious food to meet their dietary needs and food preferences for an active and healthy life (FAO, 1996). Missing meals and changing diet is a common response to food insecurity, and the latter may persist over extended periods, leading to adverse health effects, especially in children (Seligman et al., 2010). Food insecurity can result in an increased risk of death or illness from stunting, wasting, weakened responses to infection, diabetes, cardiovascular diseases, some cancers, food-borne disease and mental ill health, via insufficient quantity, poor nutritional quality of food, contaminated foods or social exclusion Friel and Ford (2015). Rising food insecurity has been strongly associated with malnutrition, sustained deterioration of mental health, inability to manage chronic disease and worse child health (Loopstra, 2014; Loopstra et al., 2015a). Food insecurity is associated with hypertension and hyperlipidaemia, cardiovascular risk factors and also with poor glycaemic control in those with diabetes, whose additional medical expenses exacerbate their food insecurity (Lee et al., 2019). Food insecurity has been found to affect school children's academic performance, weight gain and social skills (Faught et al., 2017). Obesity is more prevalent among food-insecure women, although controlling for BMI does not attenuate the association of food insecurity and chronic disease (Pan et al., 2012).

1.2 | The UK picture

The recent increase in household food insecurity in the United Kingdom is well known through the much-publicised expansion in the uptake of humanitarian aid, principally through food banks and their corresponding increase in number (Loopstra et al., 2015b). In the year ending March 2021, more than 2.5 million parcels were distributed by Trussell Trust Foodbank, and in the 6 months to March 2021, under COVID-19 lockdown measures, the number of parcels rose by 33% on the previous year (Trussell Trust, 2021).

As a nation, the United Kingdom is wealthy and one of the world's most food secure; in 2017 it was third of 113, just after Ireland and the United States (The Economist Intelligence Unit, 2019) but by April 2021 had declined to sixth place. The UK government has a legal duty under the International Covenant on Economic, Social and Cultural Rights to take appropriate steps to realise the right of everyone to be free of hunger (United Nations Office of the High Commissioner, 1966). Relative to other advanced western economies, Britain had higher general inflation, higher food, fuel and housing price inflation, lower growth in wages in the years

immediately following the 2008 global financial crisis. A letter published in the *BMJ* (Taylor-Robinson et al., 2013) alerted readers to the fact that the number of malnutrition-related admissions to hospital had doubled between 2008/9 and 2013. The United Kingdom also has a history of very large numbers of very low-paid employees; many of those accessing food banks are in work (Field et al., 2014). Persistent and widespread low pay, the proliferation of zero-hours contracts and rising living costs, especially food prices, have been suggested as contributory factors for the increase in household food insecurity (Garratt, 2015). SARS-CoV-19 pandemic arrived in the United Kingdom in early 2020 after nearly a decade of cuts to public spending, adding further pressures and uncertainty to household budgets.

For many years, the exact scale of household food insecurity in the United Kingdom was unknown. This was because there was no systematic, national assessment of the numbers of households experiencing food insecurity, but only small-scale studies (Pilgrim et al., 2012; Tingay et al., 2003). However, from 2016, the Food Standards Agency included the Adult Food Security Module of the USDA Household Food Security Survey (HFSS) (USDA, 2012) in the bi-annual Food and You Survey. The HFSS has been used in the United States and Canada to assess levels and drivers of household food security over a number of years. It contains 10 items for households without children and 18 items for households with children (age 0–17) to assess their experiences over the last 12 months. The HFSS classifies households as being food insecure when the respondent reports three or more food insecure conditions and as very low food security category if at least one member experienced reduced food intake or if insufficient resources for food disrupted eating patterns.

The latest Food and You 2 survey, Wave 1 (2020) (Armstrong et al., 2021), found that across England, Wales and Northern Ireland, 84% of respondents were classified as food secure (72% high, 12% marginal) and 16% respondents were classified as food insecure (9% low, 7% very low). Data were collected between 29th July and 6th October 2020, during the COVID-19 pandemic which had a significant societal and economic impact and key summaries are in Table 1. (For comparison, 2018 figures were 80% high food security, 10% marginally food secure, and 10% low or very low food security Table 3.) In 2020, younger adults were more likely to report that they were food insecure than older adults and food insecurity was higher on low income households. Food security was higher in most employment groups compared to those who were long term unemployed or had never worked. Food security rates were also higher in

TABLE 1 Food security England, Wales and Northern Ireland by demographic from Food and You 2 survey, Wave 1 (2020)

Demographic	Food secure
Overall 2018	90%
Overall 2020 (COVID-19 pandemic)	84%
Low income (> £ 19,000 p.a.)	68%
Professional occupations	89%
Long term unemployed	56%
Married or in a civil partnership	90%
Living as a couple	75%
Single	78%

TABLE 2 Changes of eating habits in the last 12 months for financial reasons from Food and You 2 survey, Wave 1 (2020)

Changes reported	
Respondents with children	48%
No children in the household	31%
Eaten out less	65%
Eaten at home more	66%
Cooked more at home	65%
Eaten fewer take-aways	61%
Bought items on special offer	67%
Changed where they buy food	58%
Buy cheaper alternatives	56%
Used a food bank or emergency food	5%

respondents who were married or in a civil partnership compared to those who were living as a couple or were single.

Many respondents had changed their eating habits in the last 12 months for financial reasons, as summarised in Table 2. Changing eating habits for financial reasons was more common in households with children and 5% of respondents reported that they had used a food bank or emergency food. Rising food prices can quickly lead to food insecurity with serious public health consequences (Barons & Aspinall, 2020).

1.3 | Comparison with the United States and Canada

The United States and Canada are similar to the United Kingdom in their profiles of poverty and types of government. This allows us to draw on their research where UK data and evidence is sparse (Loopstra, 2014; Tarasuk et al., 2010). Like the United Kingdom, the United States and Canada are wealthy nations with significant household food insecurity. In contrast to the United Kingdom, the United States and Canada have undertaken regular monitoring of household food security over many years through the HFSS module within regular household surveys (Tarasuk et al., 2016). This means that research on determinants and rates of food insecurity over time is more advanced and detailed in United States and Canada than in the United Kingdom. Comparative measures of poverty are shown in Table 3.

1.4 | Need for decision support

The emerging crisis in the United Kingdom is not merely a matter for charity, but of great concern to policymakers, who are legally and morally obligated to act, but may lack recent experience in dealing with needs of this kind and scale, and so require decision support. There is a need to gather what information does exist for the United Kingdom and similar countries in order to ascertain the principal drivers of household food security, and the relationships between them, to support policy-makers to design policy to tackle food security and to evaluate other policies which may impact on food security.

TABLE 3 Poverty measures across three countries

	United Kingdom	United States	Canada
Overall poverty	19.0%	11.8%	9.5%
Child poverty	26.5%	16.2%	9.0%
Working adults with no children	16.4%	–	–
Adults 18–64	–	10.7%	–
Pensioners	13.5%	9.7%	3.9%
Food security low (very low)	10.0%	11.1% (4.3%)	12.3% (2.5%)

Notes: UK absolute poverty rate measures the fraction of population with household income below 60% of (inflation-adjusted) median income in base year 2010/11. USA Census Bureau uses a set of dollar value thresholds that vary by family size and composition to determine poverty. Canada uses the Market Basket Measure, the concept of an individual or family not having enough income to afford the cost of a basket of goods and services, omitting housing and childcare costs. Food security figures are from 2018.

In ever-larger dynamic systems, such as the food security, it is increasingly difficult for decision makers (DMs) to effectively account for all the variables within the system that may influence the outcomes of interest under enactments of various policies. In particular, government policies on welfare, farming, the environment, employment, health, etc., all have an impact on food security at various levels and by various routes. Each of the influencing variables is dynamic subsystems within domain expertise, many supported by sophisticated probabilistic models. Within the food system, examples of these are medium to long range weather, which influences food supply forecast using large numerical models, and the behaviour of global markets and prices under various plausible scenarios modelled with economic models such as autoregressive or moving averages. This paper proposes an IDSS (Barons et al., 2018; Smith, 2010; Smith et al., 2015), designed to provide decision support in these types of system, for household food security in the United Kingdom. The IDSS is a computer-based tool which integrates uncertainties of different parts of a complex system and addresses the decision problem as a whole.

In Section 2, the IDSS methodology is briefly reviewed and the graphical model and inference results are presented. Section 3 details the model and variables used for utility computation in the context of food security in the United Kingdom. Then Section 4 presents the outputs and policy evaluation for the food security system. We end the paper with a discussion of our findings and the planned next steps in this research programme.

2 | INTEGRATING DECISION SUPPORT SYSTEMS

Integrating decision support systems were introduced in Smith et al. (2015, 2017), providing an unambiguous and full framework around which to evaluate the efficacy of different policy options in complex, evolving scenarios. The IDSS aids DMs in the understanding of a problem by providing a clear evaluation and comparison of the possible policy options available. It combines expert judgement with data for each subsystem resulting in a full inferential procedure able to represent complex systems. In Barons et al. (2018), we detail the iterative manner of the development of an IDSS with its DMs and expert panels. Before the elicitation starts, it is always necessary to do some preparatory work. With the help of various domain experts, the analyst will need to trawl any relevant literature and check which hypotheses on the elements of the system and relationships between them, that are found in the literature, might still be current. The analysts repeatedly review the qualitative structure of the IDSS in light of the more profound

understanding of the process acquired through more recent elicitation of the experts' causal beliefs. This modification and improvement continues until the decision centre is content that the structure is as required (Phillips, 1984). Since the process of model elicitation is an iterative one, it is often wise to begin with some simple utility measures, proceed with an initial structural model elicitation, and then revisit the initial list of attributes of the utility; detailed exploration of the science, economics or sociology can prompt the decision centre to become fully aware of the suitability of certain types of utility attribute measures. By focusing the centre and its expert panels on those issues that really impact on final outcomes we can vastly reduce the scope of a potentially enormous model; only those features that might be critical in helping to discriminate between the potential effectiveness of one candidate policy against another are required. If there is strong disagreement about whether or not a dependency exists in the system then we assume initially that a dependency does exist, except where the consensus is that its effect is weak. Further iterations of the model building process usually clarify the understanding, and if not, a sensitivity analysis can usually distinguish a meaningful inclusion from others.

The decision centre also needs to decide what time step is the most natural one to use for the purposes of the specific IDSS. This choice depends on the speed of the process, how relevant data is routinely collected on some of the components, and some technical acyclicity assumptions that are typically known only to the decision analysts. The policy decision-making cycle may also be relevant, for example, annual budget setting. There may be conflict between the granularity of contributing models of the process, sample survey regularity and the needs of the system. The granularity required is driven by the granularity of the attributes of the utility. In addition, decision analysts need to match precisely the outputs of a donating panel with the requirements of a receiving panel. When these do not naturally align, then some translation, possibly a bespoke model, may be needed between them. When expert panels design their own systems, sometimes the internal structure of one component can share variables with the internal structure of another. So, for example, flooding could disrupt both the production of food and its distribution and yet these might be forecast using different components. In such cases, the coherence of the system will be lost and the most efficient way to ensure ongoing coherence is to separate out the shared variables and ask the panels concerned to take as inputs, probability distributions from the expert panel in the shared variable, for example, flood risk.

One element of these IDSS systems is the way they can appropriately handle uncertainties associated with various modules. This is vital to reliable decision making. For example, if the inputs from one module are very speculative, and so have a high variance, then policies that work well over a wide range of such inputs will, under the sorts of risk averse decisions we have here, tend to be preferred to ones whose efficacy is very sensitive to such inputs. That is why we need conditional inputs to communicate such uncertainties.

2.1 | Technical underpinning

In this section, we briefly review the recent methodological developments to support inference for decision support as they apply here. Full details and proofs are provided in Smith et al. (2017).

Consider a vector of random variables relevant to the system $\mathbf{Y} = (Y_1, \dots, Y_n)$. Typically, there are expert panels with expertise in particular aspects of the multivariate problem. The most appropriate expert panels for each subsystem are identified, each subpanel will defer to the others, adopting their models, reasoning and evaluations as the most appropriate domain experts. Each expert panel, G_i , is responsible for a subvector \mathbf{Y}_{B_i} of \mathbf{Y} , with B_1, \dots, B_m a partition of $\{1, \dots, n\}$. The

joint model thus accommodates the diversity of information coming from the different component models and deals robustly with the intrinsic uncertainty in these submodels.

Decisions $d \in \mathcal{D}$ will be taken by a DM where \mathcal{D} represents the set of all policy options that it plans to consider. In the context of large problems like this, the DM is often a centre composed of several individuals. The DM receives information from each panel and reaches a conclusion that depends on a reward function $R(\mathbf{Y}, d)$, $\mathbf{Y} \in R_{\mathbf{Y}}$, $d \in \mathcal{D}$. Let $U(R(\mathbf{Y}, d))$ be the utility function for decision $d \in \mathcal{D}$. Our main goal is to compute the expected utilities $\{\bar{U}(d) : d \in \mathcal{D}\}$ which represents the expected utilities of a DM which is taken over the joint predictive density $f(\mathbf{y})$.

To be formally valid, any IDSS must respect a set of common knowledge assumptions shared by all panels and which comprises the union of the utility, policy and structural consensus. For a distributive IDSS, the question then becomes precisely which information each of the panels needs to donate about their areas of expertise for the maximum utility scores to be calculated. Provided that the utility function is in an appropriate polynomial form, each panel needs to deliver only a short vector of conditional moments and not entire distributions because this type of overarching framework embeds collections of conditional independence statements allowing the use of tower rule recurrences (Leonelli & Smith, 2015). This facilitates fast calculations and propagation algorithms to be embedded within the customised IDSS for timely decision-making. In such a system, individual panels can easily and quickly perform prior to posterior analyses to update the information they donate when relevant new information comes to light and this can be propagated to update the expected utility scores. This aspect of the approach is especially useful within decision support for an emergency. In any circumstances, it still represents a huge efficiency gain over having to rebuild and re-parameterise a large model. There are a number of frameworks which satisfy the requirements of the IDSS properties, including staged trees, Bayesian networks, chain graphs, multiregression dynamic models (MDM) and uncoupled dynamic Bayesian networks.

The paradigm outlined here will be illustrated throughout the remainder of the paper through a proof of concept application to an IDSS for government policy for household food security in the United Kingdom, using a MDM (Queen & Smith, 1993) as the overarching framework.

2.2 | Graphical models and Multiregression Dynamic Models

Probabilistic graphical models are particularly suited to the role of decision support as they represent the state of the world as a set of variables and model the probabilistic dependencies between the variables through a graph. In particular, the graphs can be build based on the knowledge of domain experts, provide a narrative for the system and can be transparently and coherently revised as the domain changes.

If the graph is defined as a directed acyclic graph (DAG), then the joint distribution of $\mathbf{Y} = (Y_{B_1}, Y_{B_2}, \dots, Y_{B_m})$ can be factorised as

$$f(\mathbf{y}) = \prod_{i \in [m]} f_i(y_{B_i} | y_{\Pi_{B_i}}),$$

with Π_{B_i} the indices of parents of Y_{B_i} at the panel G_i .

We consider temporal models for multivariate time series represented as a graph to account for correlation over time and within the vector of observations. Here we assume that the overall structure can be governed by an MDM and the graph is used to decompose the n -dimensional model into univariate models so that each panel G_i is composed of a subgraph and a set of

univariate nodes. It is also usually assumed that the graph structure does not change over time, that is, the dependencies between variables are static. Consider the general setting such that

$$Y_{it} \perp \mathbf{Y}_{Q_i}^t \mid \mathbf{Y}_{\Pi_i}^t, \quad \mathbf{Y}_i^{t-1}, \quad i = 1, \dots, n, \quad (1)$$

with $\{Y_{it} : i = 1, \dots, n, t = 1, \dots, T\}$ a multivariate time series composing a DAG whose vertices are univariate processes, Π_i the index parent set of Y_{it} , $\mathbf{Y}_i^t = (Y_{i1}, \dots, Y_{it})'$ the historical data and Q_i the elements of $\{1, \dots, n\}$ which are different from i and that are not in Π_i . Thus, the model assumes that each variable at time t depends on its own past series, the past series of its parents and the value of its parents at time t . This results in the joint density function

$$f(\mathbf{y}) = \prod_{t=1}^T \prod_{i=1}^n f_{i,t}(y_{it} \mid \mathbf{y}_{\Pi_i}^t, \mathbf{y}_i^{t-1}). \quad (2)$$

The temporal evolution is defined through the observation and system equations given by

$$\begin{aligned} Y_{it} &= \mathbf{F}'_{it} \boldsymbol{\theta}_{it} + \epsilon_{it}, \\ \boldsymbol{\theta}_{it} &= G_{it} \boldsymbol{\theta}_{i,t-1} + \boldsymbol{\omega}_{it}, \end{aligned}$$

with $\epsilon_{it} \sim N[0, V_{it}]$ and $\boldsymbol{\omega}_{it} \sim N[\mathbf{0}_d, W_{it}]$ and $\boldsymbol{\theta}_{it} \in \Theta_i \subset \mathfrak{R}^d$. The errors are assumed to be independent of each other and through time and \mathbf{F}_{it}, G_{it} are assumed to be known at time t . Given the initial information, $\boldsymbol{\theta}_{i0} \mid \mathcal{I}_0 \sim N[\mathbf{m}_{i0}, C_{i0}]$, the parameters $\boldsymbol{\theta}_{it}$, $i = 1, \dots, n$ may be updated independently given the observations at time t . Conditional forecasts may also be obtained independently. These results are proved in Queen and Smith (1993) assuming Gaussian distributions for the error terms. The predictive density is given by

$$f_{i,t}(y_{it} \mid \mathbf{y}_{\Pi_i}^t, \mathbf{y}_i^{t-1}) = \int_{\Theta_i} g_{it}(y_{it} \mid \mathbf{y}_{\Pi_i}^t, \mathbf{y}_i^{t-1}, \boldsymbol{\theta}_{it}) \pi_i(\boldsymbol{\theta}_{it} \mid \mathbf{y}_{\Pi_i}^{t-1}, \mathbf{y}_i^{t-1}) d\boldsymbol{\theta}_{it}. \quad (3)$$

Let $\mathbf{D}_t = (\mathbf{y}_t, \mathbf{D}_{t-1})$ be the information available at time t . Inference about $\boldsymbol{\theta}_{it}$ is based on forward filtering equations to obtain posterior moments at time t :

- Posterior distribution at time $t - 1$: $\boldsymbol{\theta}_{i,t-1} \mid \mathbf{D}_{t-1} \sim N[\mathbf{m}_{i,t-1}, C_{i,t-1}]$;
- Prior distribution at time t : $\boldsymbol{\theta}_{it} \mid \mathbf{D}_{t-1} \sim N[\mathbf{a}_{it}, R_{it}]$, with $\mathbf{a}_{it} = G_{it} \mathbf{m}_{i,t-1}$ and $R_{it} = G_{it} C_{i,t-1} G'_{it} + W_{it}$;
- One step ahead prediction: $\mathbf{y}_{it} \mid \mathbf{y}_{\Pi_i}^t, \mathbf{D}_{t-1} \sim N[\mathbf{f}_{it}, Q_{it}]$, with $\mathbf{f}_{it} = \mathbf{F}'_{it} \mathbf{a}_{it}$ and $Q_{it} = \mathbf{F}'_{it} R_{it} \mathbf{F}_{it} + V_{it}$;
- Posterior distribution at time t : $\boldsymbol{\theta}_{it} \mid \mathbf{D}_t \sim N[\mathbf{m}_{it}, C_{it}]$, with $\mathbf{m}_{it} = \mathbf{a}_{it} + A_{it} \mathbf{e}_{it}$ and $C_{it} = R_{it} - A_{it} Q_{it} A'_{it}$ and $\mathbf{e}_{it} = \mathbf{y}_{it} - \mathbf{f}_{it}$, $A_{it} = R_{it} \mathbf{F}_{it} Q_{it}^{-1}$.

If data are observed from time 1 to T , then backward smoothing may be used to obtain the posterior moments of $\boldsymbol{\theta}_{it} \mid \mathbf{D}_T$, $t = 1, \dots, T$. Thus,

$$\boldsymbol{\theta}_{it} \mid \boldsymbol{\theta}_{i,t+1}, \mathbf{D}_T \sim N(\mathbf{h}_{it}, H_{it}),$$

with $\mathbf{h}_{it} = \mathbf{m}_{it} + C_{it} G'_{i,t+1} R_{i,t+1}^{-1} (\boldsymbol{\theta}_{i,t+1} - \mathbf{a}_{i,t+1})$, $H_{it} = C_{it} - C_{it} G'_{i,t+1} R_{i,t+1}^{-1} G_{i,t+1} C_{it}$ and $\mathbf{h}_{iT} = \mathbf{m}_{iT}$ e $H_{iT} = C_{iT}$ the initial values.

Modelling via dynamical models depends on specifying the variances, V_{it} and W_{it} . For the state variance W_{it} , instead of estimation we adopt the idea of discount factors (West & Harrison, 1997).

In particular, $W_{it} = C_{it}(1 - \delta_i)/\delta_i$, $\delta_i \in (0, 1)$ which for each t implies $R_{it} = C_{it}/\delta_i$, representing the loss of information over time. Note that δ_i close to 1 indicates stability while δ_i small indicates larger variance over time and, consequently, larger loss of information from $t - 1$ to t . Typically δ_i is defined to be between 0.7 and 1, with larger values representing smoother processes.

The MDM framework is extended to allow variances to vary stochastically over time, analogously to the mean parameters. Thus, embracing the approach described in West and Harrison (1997), we let the variance $V_{it} = \phi_{it}^{-1}$ and $\phi_{i,t-1} | D_{t-1} \sim \text{Gamma}(n_{i,t-1}/2, d_{i,t-1}/2)$. The Gamma evolution model is given by

$$\phi_{it} | D_{t-1} \sim \text{Gamma}(\delta_i^* n_{i,t-1}/2, \delta_i^* d_{i,t-1}/2), \quad (4)$$

with $\delta_i^* \in (0, 1)$ being the discount factors. For this evolution $E[\phi_{it} | D_{t-1}] = E[\phi_{i,t-1} | D_{t-1}]$ but $\text{Var}[\phi_{it} | D_{t-1}] = \text{Var}[\phi_{i,t-1} | D_{t-1}]/\delta_i^*$ implying an increase in the variance as time evolves controlled by δ_i^* . The posterior distribution at time t is obtained analytically as $\phi_{it} | D_t \sim \text{Gamma}(n_{it}/2, d_{it}/2)$ with $n_{it} = \delta_i^* n_{i,t-1} + 1$ and $d_{it} = \delta_i^* d_{i,t-1} + S_{i,t-1} e_{it}' Q_{it}^{-1} e_{it}$, with $S_{i,t-1} = d_{i,t-1}/n_{i,t-1}$. This conjugacy results in closed-form recurrence updating equations for this variance model.

The discount factors (δ_i, δ_i^*) , $i = 1, \dots, n$ need to be specified in the FFBS algorithm. To preserve computational simplicity, we follow the grid search approach as used in Costa et al. (2019) and select the best configuration of (δ_i, δ_i^*) using model comparison via Bayes factors (Kass & Raftery, 1995). In particular, the marginal likelihood was approximated using the Shifted-Gamma estimator (Raftery et al., 2007). Given simulations from the analytical posterior distributions of (θ_t, ϕ_t) , the density in the observational equation $p(\mathbf{y}_t | \theta, \phi, \delta, \delta^*)$ may be evaluated for all simulated state parameters and the densities may be used to estimate the marginal distribution of \mathbf{Y}_t given (δ, δ^*) . This can be easily repeated for several competing models and the larger Bayes factor indicates the best model. Note that, the discount factors will also control for overfitting by selecting values closer to 1 when the stochastic evolution of mean process is smoother over time, in the limit a constant regression coefficient could be obtained for $\delta_i = 1$ and a constant variance for $\delta_i^* = 1$. Here the uncertainty about the discount factors is considered through the Bayes Factor (Kass & Raftery, 1995). This measures the evidence provided by data in favour of model M_1 compared to model M_2 and is given by

$$B_{12} = \frac{f(\mathbf{y} | \delta_1, \delta_1^*)}{f(\mathbf{y} | \delta_2, \delta_2^*)}. \quad (5)$$

The predictive densities $f(\mathbf{y} | \delta, \delta^*)$ are approximated using the shifted-gamma estimator (Raftery et al., 2007) which considers the sequence of log-likelihood values $\{l_k = \log(f(\mathbf{y} | \theta^{(k)}, \delta, \delta^*)): k = 1, \dots, M\}$ with $\theta^{(k)}$ simulated from its posterior distribution. In the model choice problem, $2 \log(B_{12})$ greater than 2 indicates positive evidence in favour of Model 1 and values greater than 10 indicates very strong evidence according to guidelines in Kass and Raftery (1995).

2.3 | Expected utility computation and scenario evaluation

The predictive posterior distribution for a replicated observation \tilde{y} is obtained using $f(\tilde{\mathbf{y}})$ as defined in Equations (2) and (3). When the utility function is assumed to be linear, then $U(R(\mathbf{Y}, d)) = \sum_{i \in [n]} k_i U_i(R_i(\mathbf{Y}_i, d))$, so that the expected utility is given by

$$\bar{U}(d) = \sum_{i \in [n]} k_i \bar{U}_i(d | y_{\Pi_i}), \text{ and}$$

$$\bar{U}_i(d | y_{\Pi_i}) = \int_{\Theta_i} \int_{\mathcal{X}_{y_i}} U_i(R_i(y_{it}, d)) g_{it}(y_{it} | y_{\Pi_i}^t, y_i^{t-1}, \theta_i) \pi_i(\theta_{it} | y_{\Pi_i}^t, y_i^{t-1}) dy_i d\theta_i.$$

If $U(R_i(\tilde{y}_{it}, d))$ are linear functions of \tilde{y}_{it} , then the expected utilities can be computed analytically using chain rules of conditional probabilities. If $U(R_i(\tilde{y}_{it}, d))$ is a nonlinear function of \tilde{y}_{it} then expected values are computed by Monte Carlo integration (Robert & Casella, 2004). Suppose that $\theta_{1:T}$ was simulated using the forward filtering and backwards sampling algorithm as described in Section 2.2. Then, $\bar{U}_i(d | y_{\Pi_i})$ can be obtained by simulating from the observation density $g_{it}(\cdot | \tilde{y}_{\Pi_i}^t, \tilde{y}_i^{t-1}, \theta_{it})$. Note that some ordering in computing expectations need to be followed, starting from the variables such that $\mathcal{L}_i(\mathbf{Y}_{it}) = \emptyset$, their descendants and so on.

The types of overarching descriptions suitable for these applications must be rich enough to explore both the effects of shocks to the system and the application of policies. These can be conveniently modelled through chains of causal relationships, where causal means that there is an implicit partial order to the objects in the system and we assume that the joint distributions of variables not downstream of a controlled variable remain unaffected by that control. The downstream variables are affected in response to a controlled variable in the same way as if the controlled variable had simply taken that value. Note that in the case the controlled variables were simulated from their posterior predictive distributions then the downstream variables will also be simulated conditional on each value.

3 | IDSS: UK FOOD SECURITY

This section presents the application of the framework described in Section 2.2 to the context of UK food security. The nodes, panels, utility function and graph structure used for this application are all described here.

Following a literature search to identify the key issues surrounding household-level food security in the United Kingdom (summarised in Sections 1.1–1.3), a series of decision conferences was held with Warwickshire County Council and other local public services. Since food security is not a discrete responsibility of any one local authority department, delegates attended from the council's public health, legal & governance, Warwickshire Observatory (data and statistics), corporate GIS, renewable energy, social & financial inclusion, localities & partnerships, child poverty, education, emergency planning, libraries & customer services, and corporate policy departments. The events were also attended by representatives from Warwickshire Police.

These delegates were engaged in what can be called joint model building or soft elicitation (French, 2021; Wilkerson & Smith, 2021). First they were presented with the UK picture and then asked to express their beliefs about how this related specifically to Warwickshire. The academics formulated the experts' beliefs into a probabilistic graphical model. Over several workshops, the semantics of the experts' beliefs on the structure of the system were clarified and a consensus model was produced which reflected these beliefs. Part of the process was to discuss what were the relevant granularities of data needed to support decision-making adequately while maintaining a model as parsimonious as possible. Since Food Security was not measured, these experts gave advice on what proxy measures would be suitable.

3.1 | Structure of the IDSS

For potentially massive and very heterogeneously informed graphical models, it is usually wise to elicit the graphical framework directly from experts who understand the interdependences between components of the system rather than relying on automatic selection methods. As well as sidestepping the difficult technical issue surrounding the choice of an appropriate score function for models, it also ensures that the structural framework around which inferences take place is meaningful and defensible in the decision-makers' context. Furthermore, complete data sets across the whole composite are rare. Therefore, the structure of this graphical model with the variables influencing food security was elicited from the experts. In this application, instead of doing model comparison for selecting the best network structure we rely on expert elicitation for the topology. In this case, expert opinion guides the choice of nodes and links between nodes. This choice aims to maintain the causal perspective in the graph allowing for cause-effect inferences as a result. For further details on how to elicit the graph structure manually based on domain expert information, see Smith (2010), Kjaerulff and Madsen (2013) and Barons et al. (2018).

We first elicited the main variables of interest, then the variables which affect those variables and so on until a suitable level of detail has been obtained. This was effected using an iterative process, drawing on the food poverty literature and checking with domain experts, refining and repeating. In particular, the general framework was confirmed by work produced independently in Loopstra (2014). The variables and their dependencies for the UK food system are shown in Figure 1, which illustrates the 16-node graph structure obtained through literature and confirmed by the experts.

The interaction with policymakers has ensured we have the required structural, utility and policy consensus to make up the CK-class required by the IDSS and we check that the conditions needed provide sufficient information to fully and unambiguously define the composite probability model of the whole process (for details see Smith et al., 2017).

3.2 | Expert panels

Having identified the factors influencing household food security in the United Kingdom, the next step is to identify the most relevant experts to provide data on these. The panels constituted for such an IDSS will often be chosen to mirror the panels that are already constituted for similar purposes, for example, in the United Kingdom, the Office for Budget Responsibility, HM Treasury and The Confederation of British Industry all produce economic forecasts on the UK Economy. Looking at where the relevant information is held gives some very natural panels. See Appendix C (p20) for details of data sources.

The 16-node graph structure illustrated in Figure 1 becomes a 9-panel IDSS (Figure 2) as sources of data on the variables are matched with holders and experts and their models concerning that data, meaning some nodes merge into the same panel. Panel G2 reports on cost of food given inputs from panel G5 on food supply, incorporating both the variables food imports and domestic food production in Figure 1. Panel G5, in turn, relies on information from G8, the Met office, on weather and climate patterns to calculate its expectations of food supply, since both domestic and world production and supply chain disruption are weather related. Household income, G1, impacts directly on the utility. Panel G1 relies on information provided by G3 (incorporating access to credit, benefits and tax) and G4 (incorporating cost of housing and energy) to make its predictions under different policy scenarios. G4 advises on cost of living including

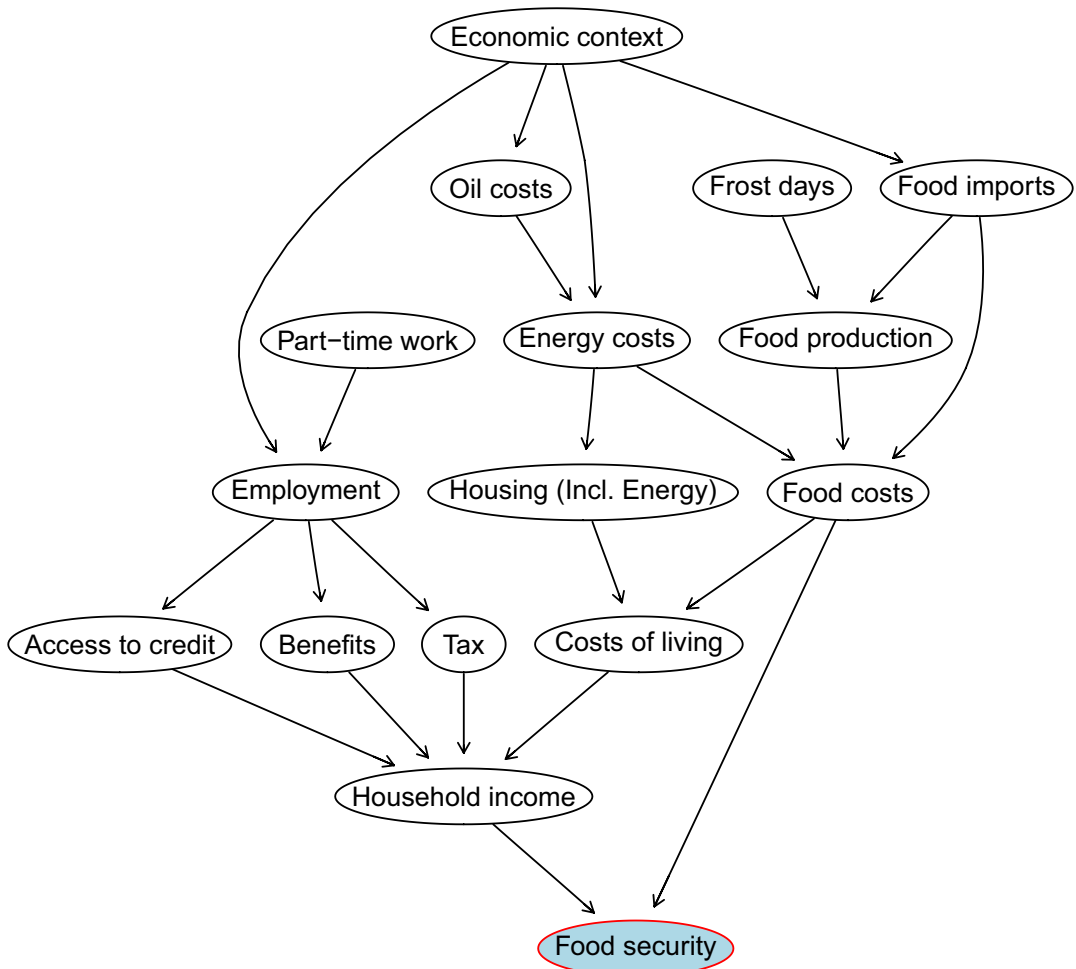


FIGURE 1 Proposed graph structure and nodes for UK food security decision support

energy, housing and other essentials. G3 assesses income taking into account employment, tax and social security, taking inputs from G7 and G9. G7 advises on demography, including single parents, immigrants, disability and those with no recourse to public funds. G9 advises on matters of the economy and informs the oil price panel, G6, and the cost of living panel, G4 as well as G3. Each panel provides summaries of their model outputs for each of the policy decisions under consideration, to the panels downstream, which condition their models on those summaries and the policy under consideration.

3.3 | MDM IDSS for food security

In every decision support scenario, it is essential to clarify the goals of the DM. Support for household food security is provided in the UK context through local government, typically city or county councils through their financial inclusion and child poverty policies. City, county or district councils in the United Kingdom fulfil their statutory obligations to meet the requirements of central government. However, in addition, they go beyond mere compliance to represent and

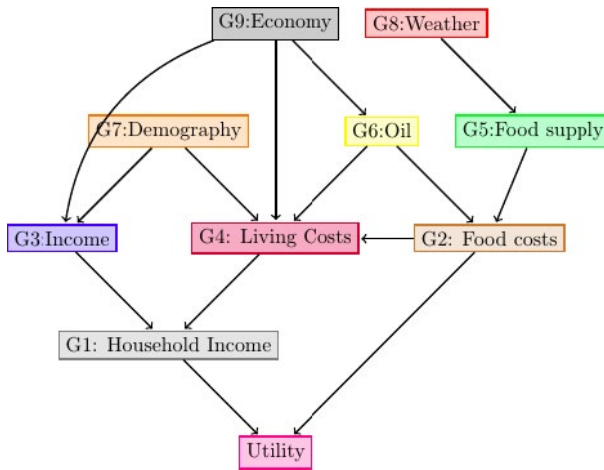


FIGURE 2 The expert panels required for this integrating decision support system. Each node represents an expert panel which, using its models and data, provides summaries of expected values and relevant moments under each policy decision being considered. Some panels use models subsuming several of the items in Figure 1. See Section 3.3 for details

reflect their local communities and continually improve the lives of the citizens within their geographical region, with a special focus on improving the circumstances of the most disadvantaged. In this context, policy and scenario comparison through an IDSS can explicitly present directions for improvements in food security in the United Kingdom.

Here we assume plausible models for the expert panels and utility, based on publicly available data. The attributes being measured to compose the food model were obtained at the Office for National Statistics which publishes official statistics for the United Kingdom. The time series for all nodes are measured annually and the temporal window considered goes from 2008 to 2018. Each variable is detailed in Appendix A.

3.3.1 | Utility function elicitation

In order to construct this IDSS for food security, we defined the utility function and developed a suitable mathematical form for it. One candidate measure of household food security was data from food bank charities. However, studies have shown that food bank use is not a good measure of food poverty (Coleman-Jensen et al., 2016; Kirkpatrick & Tarasuk, 2009). In the absence of a direct measure of household food security in the region, we consulted Warwickshire County Council, who identified education, health, cost and social unrest as suitable attributes for a utility.

In constructing a utility function based on these attributes, it appeared appropriate to assume value independence (Keeney & Raiffa, 1993). Let Z_1 = measures of education, Z_2 = measures of health, Z_3 = Measures of social unrest, Z_4 = cost of ameliorating policies to be enacted. We then specified suitable forms for the marginal utility functions. For social unrest, health and education was modelled as exponential, while the utility on cost was modelled as linear. It was therefore decided that one family of appropriate utility functions might take the form:

$$U(z) = a + bz_4 + \sum_{i=1}^3 1 - \exp(-c_i z_i), \tag{6}$$

where $z = (z_1, z_2, z_3, z_4)$ and whose parameters (a, b, c_1, c_2, c_3) were then elicited. Note that social unrest was omitted in this analysis since, at this stage, we had not performed the underlying necessary elicitation sessions. The cost of the policy is not explicit in the utility calculated here; the decision-makers will consider costs alongside the utility scores for candidate policies assessed by the IDSS.

For the purposes of this proof of concept, health and education indicators were considered as proxies of food security at the household level. One requirement of the attributes of a utility function is that they must be measurable; it must be possible to say whether an event has happened or a threshold has been reached. The Health and Education indicators defined below satisfy these requirements:

- Health—the count of finished admission episodes with a primary or secondary diagnosis of malnutrition coded ICD-10. A ICD-10 code of malnutrition on the episode indicates that the patient was diagnosed with, and would therefore being treated for malnutrition during the episode of care.
- Education—the gap index measuring the differences between the disadvantaged and non-disadvantaged groups in key stages 2 and 4 (Hill, 2014). The index is the mean rank for all the disadvantaged and non-disadvantaged pupils divided by the number of pupils in each cohort. This decimal mean rank difference is scaled to 10 and ranges from 0 to 10, where a higher value means a higher attainment of non-disadvantaged compared to disadvantaged pupils. The index aims to be resilient to changes in the grading systems and in the assessments and curricula, and may be used for temporal comparisons.

3.3.2 | Panel models

Health and education are directly affected by household income (HIncome, panel G_1) and food costs (CFood, panel G_2). The variables are modelled in the log scale as both are percentages or rate:

$$\begin{aligned} \log(\text{Health}_t) &= \delta_{01,t} + \delta_{11,t} \text{HIncome}_t + \delta_{21,t} \text{CFood}_t + \epsilon_{ht}, \\ \log(\text{Education}_t) &= \delta_{02,t} + \delta_{12,t} \text{HIncome}_t + \delta_{22,t} \text{CFood}_t + \epsilon_{et}. \\ (\text{Health}_t) &= \delta_{01,t} + \delta_{11,t} \text{HIncome}_t + \delta_{21,t} \text{CFood}_t + \epsilon_{ht}, \\ (\text{Education}_t) &= \delta_{02,t} + \delta_{12,t} \text{HIncome}_t + \delta_{22,t} \text{CFood}_t + \epsilon_{et}. \end{aligned} \quad (7)$$

Panel G_1 advises on disposable household income after accounting for the cost of living (panel G_4), taxes and also the access to credit and benefits (panel G_3).

$$\text{HIncome}_t = \theta_{01,t} + \theta_{11,t} \text{Lending}_t + \theta_{21,t} \text{Tax}_t + \theta_{31,t} \text{Benefits}_t + \theta_{41,t} \text{CLiving}_t + \epsilon_{1t}.$$

The variable costs of food (Panel G_2) depends on costs of energy (panel G_6) and on food supply, imports and exports and food production (panel G_5).

$$\text{CFood}_t = \theta_{02,t} + \theta_{12,t} \text{FProduction}_t + \theta_{22,t} \text{FImports}_t + \theta_{32,t} \text{CEnergy}_t + \epsilon_{2t}.$$

Panel G_3 reports on variables affecting the income such as lending, tax and unemployment. Unemployment depends on the economic context (panel G_9) represented by GDP and on part-time workers (panel G_7).

$$\begin{aligned} \text{Lending}_t &= \theta_{03,t} + \theta_{13,t} \text{Unemployment}_t + \epsilon_{3t}, \\ \text{Tax}_t &= \theta_{03,t}^* + \theta_{13,t}^* \text{Unemployment}_t + \epsilon_{3t}^*, \\ \text{Benefits}_t &= \theta_{03,t}^{**} + \theta_{13,t}^{**} \text{Unemployment}_t + \epsilon_{3t}^{**}, \\ \text{Unemployment}_t &= \theta_{03,t}^{***} + \theta_{13,t}^{***} \text{Part-time}_t + \theta_{23,t}^{***} \text{GDP}_t + \epsilon_{3t}^{***}. \end{aligned}$$

Panel G_4 reports on costs of living which depend on costs of food (panel G_2), on costs of housing including energy. Costs of housing depend on costs of energy (panel G_6).

$$\begin{aligned} \text{Cliving}_t &= \theta_{04,t} + \theta_{14,t} \text{CFood}_t + \theta_{24,t} \text{CHousing}_t + \epsilon_{4t}, \\ \text{CHousing}_t &= \theta_{04,t}^* + \theta_{14,t}^* \text{CEnergy}_t + \epsilon_{4t}^*. \end{aligned}$$

Panel G_5 (Food supply) reports on food production and imports which depend on the economic context (panel G_9):

$$\begin{aligned} \text{FProduction}_t &= \theta_{05,t} + \theta_{15,t} \text{GDP}_t + \theta_{25,t} \text{FImports}_t + \epsilon_{5t}, \\ \text{FImports}_t &= \theta_{05,t}^* + \theta_{25,t}^* \text{GDP}_t + \epsilon_{5t}^*. \end{aligned}$$

Panel G_6 reports on oil costs and energy given inputs from panel G_9 about economic context.

$$\begin{aligned} \text{COil}_t &= \theta_{06,t} + \theta_{16,t} \text{GDP}_t + \epsilon_{6t}, \\ \text{CEnergy}_t &= \theta_{05,t}^* + \theta_{15,t}^* \text{COil}_t + \epsilon_{5t}^*. \end{aligned}$$

Panels G_7 (Demography), G_8 (Weather) and G_9 (Economy) report on demography, weather and economic context, respectively, with model equations given by

$$\begin{aligned} \log(\text{PartTime}_t) &= \theta_{07,t} + \epsilon_{7t}, \\ \text{Frost}_t &= \theta_{08,t} + \epsilon_{8t}, \\ \text{GDP}_t &= \theta_{09,t} + \epsilon_{9t}. \end{aligned}$$

Using these models as the panels' models to obtain the predictive distribution in Equations (2) and (3), we now define the utility function used to compare a number of scenarios.

4 | MODEL OUTPUTS AND SCENARIO EVALUATION

The MDM dynamic coefficients and variances were estimated based on the best hyperparameter configuration as detailed in Section 2.2. With 11 time points, we keep the model parsimonious and consider the discounts factors $\delta_i = \delta_0$ and $\delta_i^* = \delta_0^*$, $i = 1, \dots, n$ for the evolutions in the mean and variance models, respectively. Table 4 presents the model comparison for a grid of values for (δ_0, δ_0^*) . The best model has $(\delta_0, \delta_0^*) = (0.85, 0.95)$ and these values were used to evaluate the posterior distribution of the utility function of interest.

TABLE 4 Log Bayes factor for model M_1 with discount (δ_0, δ_0^*) versus model M_2 with discount $(0.95, 0.90)$

		δ_0^*		
		0.90	0.95	0.99
δ_0	0.85	66.24	71.14	70.75
	0.90	35.61	41.48	37.29
	0.95	22.62	13.39	17.69
	0.95	3.08	0.00	0.06

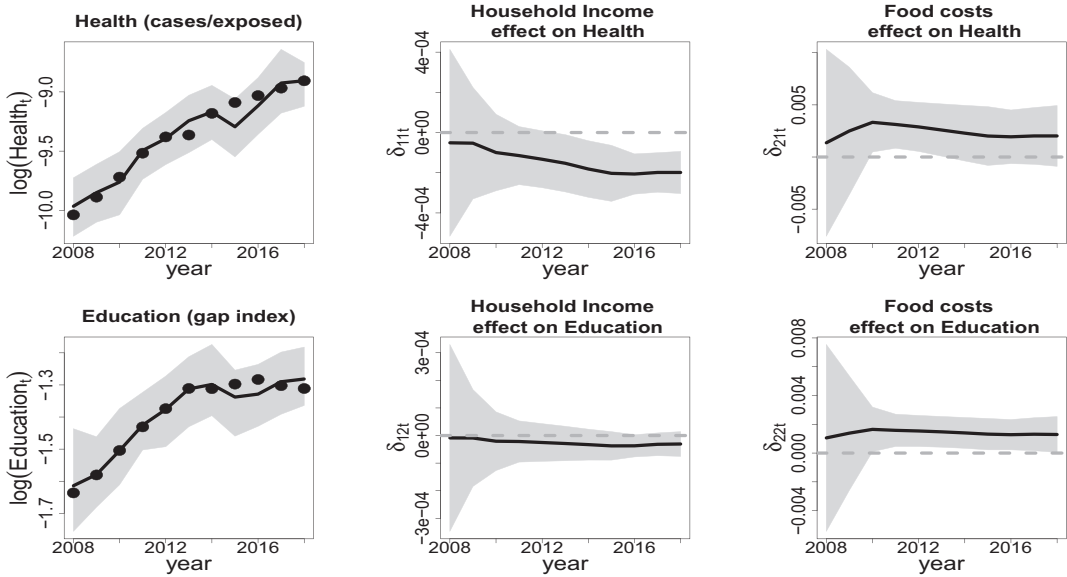


FIGURE 3 Attributes composing the utility function, effects of household income and food costs and MDM fit (mean and 95% credible interval), 2008–2018

Figure 3 presents the fit and effects of household income and food costs on health and education obtained by recursively updating of posterior moments based on the forward filtering and backward algorithm presented in Section 2.2. All the relations indicated by the experts in the field are verified by the dataset gathered for this application. Notice the negative effect of household income and positive effect of food costs on the rate of malnutrition. The effect of household income on education is mostly not significant or negative over the observed temporal window and the effect of food costs is mostly positive on the percentage of disadvantaged pupils. Figure 4 presents the MDM fit for all the variables in the food security network. Note that uncertainty is well captured by the posterior intervals. Frost days presents the largest uncertainty, indicating that more granularity could be used in this panel of experts to better explain weather effects. The dynamic coefficient effects (not shown here) are not constant over time, indicating that our proposal adequately accounts for non-stationarities.

After fitting the dynamical model, four different policies were compared using the IDSS approach described in Section 2.

Policy 1 is ‘do nothing’, that is, all variables kept at the baseline observed values. Policy 2 accounts for an increase of 25% in food costs driven by economic or political policy, such as

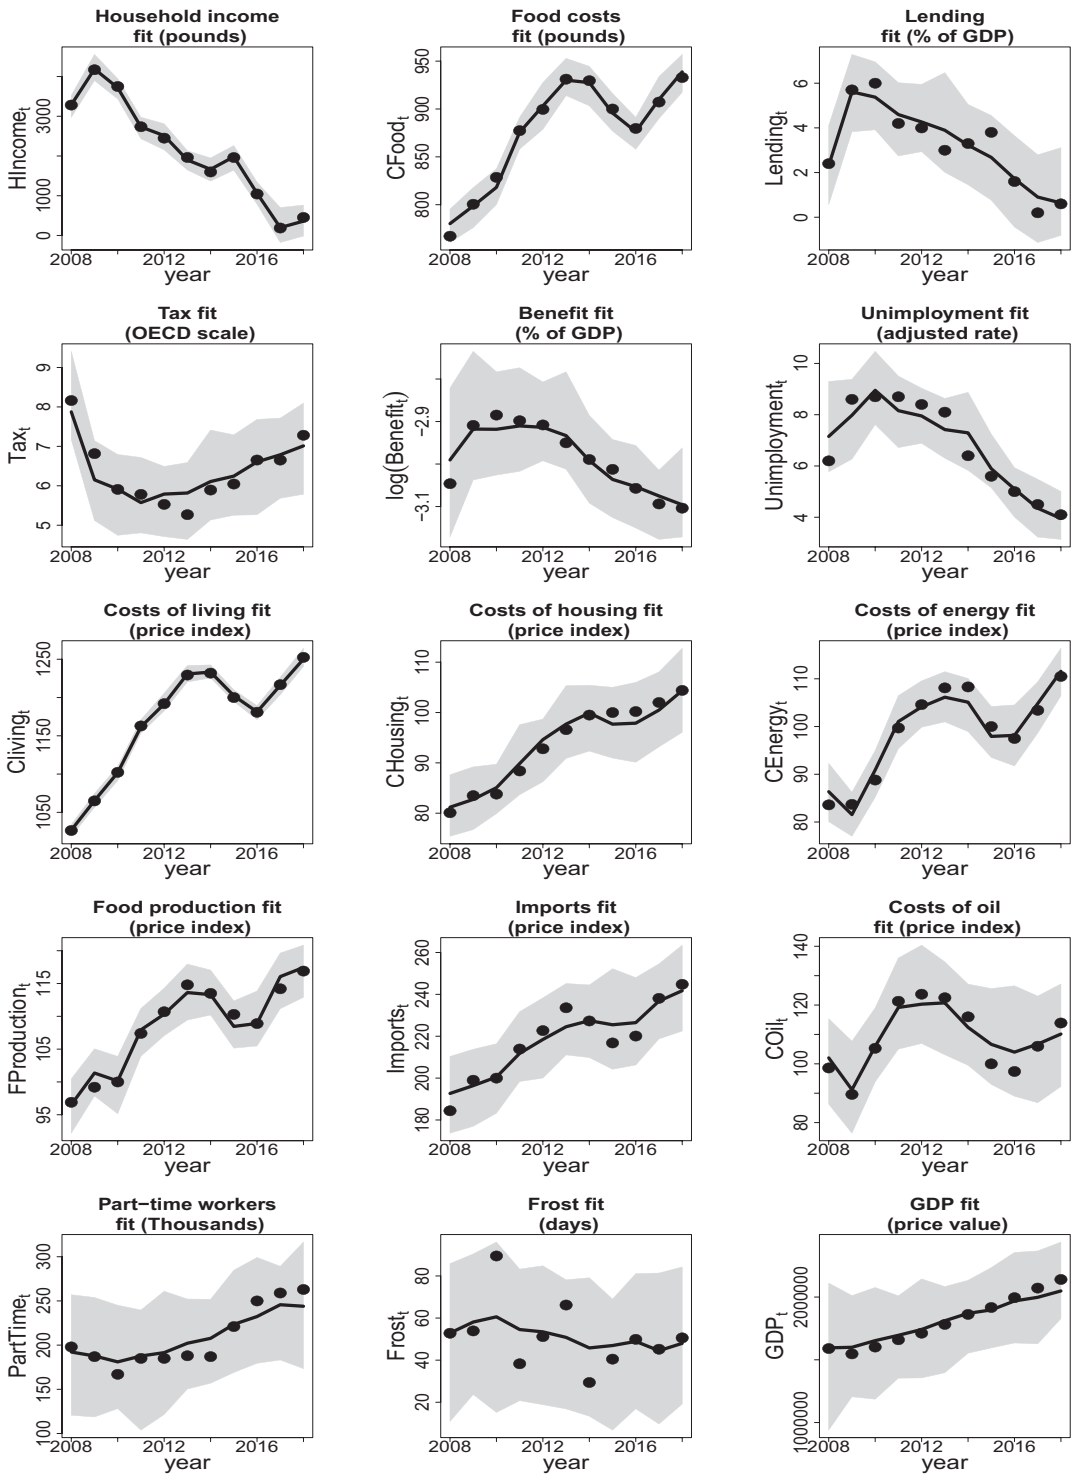


FIGURE 4 Variables composing the food network and dynamical regression model fit (mean and 95% credible interval), 2008–2018

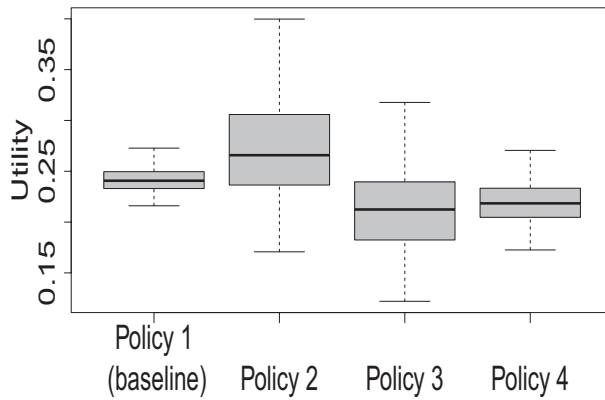


FIGURE 5 Boxplot of utility function posterior distribution

Brexit (Barons & Aspinall, 2020). Policy 3 represents a subsidy policy leading to a decrease of 25% in food costs. Policy 4 is a compound economic, welfare and incentive policy leading to a 15% reduction in food prices plus an increase in household income by 15%. The expected value of utility for policies 1, 2, 3 and 4 are 0.2400, 0.2808, 0.2091 and 0.2232, respectively. Small values for the utility are associated with smaller rates of malnutrition and smaller percentage of disadvantaged pupils.

We see that Policy 3, a 25% decrease in food costs, gives the lowest (best) utility score. This scores better than policy 4 which decreases food costs by 15% while raising incomes by 15%. Policies leading to an increase in food prices of 25% are clearly worse than baseline, as expected.

Figure 5 presents the posterior utility function for the 4 policies. The baseline policy presents the smaller spread reflecting the smaller uncertainty in this scenario.

Different representations of the utility outputs are suitable for different actors within the decision-making process. In the local council example, council officers have expertise in disparate domains and often make recommendations on courses of action within their remit to the elected members who make the final decisions. If elected members raise queries, decisions and recommendations can undergo further scrutiny before a final decision is made. If the IDSS is made to support council officers with relevant expertise, the plots in Figure 5 might prove useful, as the entire distribution is shown. However, for decision-makers with less technical expertise, a simpler representation, giving utility score along with a natural language output might be more suitable, as discussed in (Barons et al., 2018). For instance, Figure 6 presents the posterior median for the utility function for each policy divided by the baseline median.

5 | DISCUSSION AND FURTHER DEVELOPMENTS

Bespoke decision support based on the IDSS paradigm is increasingly being developed in disparate domains. Here we have shown a proof of concept IDSS for policymakers concerned with ameliorating household food security in the United Kingdom. We have identified the main drivers of UK food security, drawing partly on research from the United States and Canada where food security has been measured for a number of years and therefore the understanding of

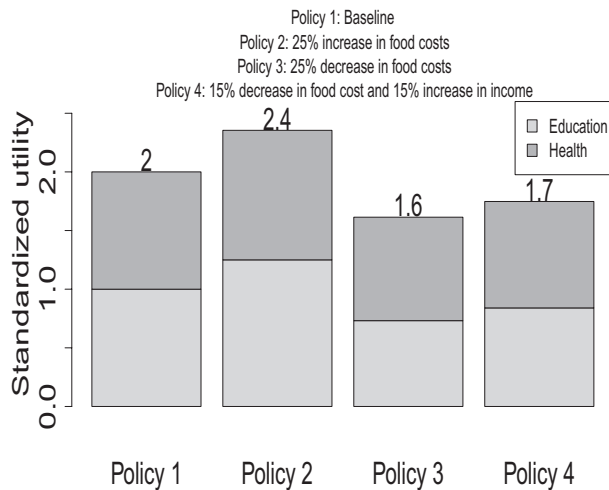


FIGURE 6 Standardised utility posterior median decomposed for health and education variables

determinants of household food security is more advanced than in the United Kingdom. We have identified plausible expert panels based on UK structures to provide inputs for the IDSS and have constructed models based on publicly available data.

For this particular application to food insecurity in the United Kingdom, dynamical models were fitted to the time series and a good fit was obtained for the selected variables. An alternative would be to consider economic models which would rely on the input from experts in the field.

Furthermore, we have demonstrated the output of the IDSS under a number of plausible policies by computing the posterior distribution of the proposed utility function. In particular, we have assumed equal weighting between health and educational attainment as a proxy for food insecurity in the UK local authority setting. The posterior distribution obtained for several policies may be compared in terms of spread as well as location measures.

To move from a proof of concept to a working IDSS, one task would be to elicit the user preferences in displaying the results, as discussed in Barons et al. (2018, 2021).

We have found the IDSS framework really useful to elicit the vital features of a problem and their relationships. Using the IDSS ensures that these domain judgements derive from those panels of experts that understand the nuances of a particular field so that evaluations of different policy options reflect these expert judgements and compose them in a logical way. The fact that the composite explicitly and transparently is framed by a formal graph means that its outputs can be understood and their genesis from the composite explained. They can therefore be intelligently discussed and if necessary modified through both better modelling and better information about the key attributes driving the decisions. With regard to the latter we can expect that as data from surveys like Food and You are progressively collected, these can enhance the models we discuss above enabling the IDSS to better evaluate the efficacy of different options making an IDSS analysis, like the one above, even more discriminating and helpful to policy makers.

ACKNOWLEDGEMENT

This research was supported by EPSRC grant EP/K039628/1.

ORCID

Martine J. Barons  <https://orcid.org/0000-0003-1483-2943>

Thais C. O. Fonseca  <https://orcid.org/0000-0002-4943-3259>

REFERENCES

- Armstrong, B., King, L., Clifford, R. & Jitlal, M. (2021) Food and you 2: Wave 1 key findings. *Technical report*, Food Standards Agency UK.
- Barons, M.J. & Aspinall, W. (2020) Anticipated impacts of Brexit scenarios on UK food prices and implications for policies on poverty and health: a structured expert judgement approach. *BMJ Open*, 10, e032376.
- Barons, M.J., Wright, S.K. & Smith, J.Q. (2018) Eliciting probabilistic judgements for integrating decision support systems. In: Dias, L.C., Morton, A. & Quigley, J. (Eds.) *Elicitation*. Cham, Switzerland: Springer, pp. 445–478.
- Barons, M., Bhatia, S., Double, J., Fonseca, T., Green, A., Krol, S. et al. (2021) Safeguarding the nation's digital memory: towards a Bayesian model of digital preservation risk. *Archives and Records*, 42, 58–78.
- Coleman-Jensen, A., Rabbitt, M., Gregory, C. & Singh, A. (2016) Household food security in the United States in 2015. Economic research report no. (ERR-215) September 2016. *Technical report*, USDA.
- Costa, L., Smith, J.Q. & Nichols, T. (2019) A group analysis using the Multiregression Dynamic Models of fMRI networked time series. *Journal of Statistical Planning and Inference*, 198, 43–61.
- FAO. (1996) Rome declaration on world food security. *Technical report*, Food and Agriculture Organization of the United Nations.
- Faught, E., Williams, P., Willows, N., Asbridge, M. & Veugelers, P. (2017) The association between food insecurity and academic achievement in Canadian school-aged children. *Public Health Nutrition*, 20, 1–8.
- Field, F., Thornton, T., Glen, J., Jenkin, B., Lewell-Buck, E. & Newton, S. (2014) Feeding Britain: a strategy for zero hunger in England, Wales, Scotland and Northern Ireland, the report of the all-party parliamentary inquiry into hunger in the United Kingdom. *Technical report*, All-Party Parliamentary Inquiry into Hunger in the United Kingdom.
- French, S. (2021) From soft to hard elicitation. *Journal of the Operational Research Society*, 1–17. <https://doi.org/10.1080/01605682.2021.1907244>
- Friel, S. & Ford, L. (2015) Systems, food security and human health. *Food Security*, 7, 437–451. <https://doi.org/10.1007/s12571-015-0433-1>
- Garratt, E. (2015) Csi 13: food insecurity and foodbank use. *Technical report*, Centre for Social Investigations, Oxford.
- Hill, K. (2014) Measuring disadvantaged pupils attainment gaps over time: working methodology and statistics for 2012–2014. Working paper SFR 40/2014, Department for Education.
- Kass, R. & Raftery, A.E. (1995) Bayes factor. *Journal of the American Statistical Association*, 90, 773–795.
- Keeney, R.L. & Raiffa, H. (1993) *Decision with multiple objectives: preferences and value trade-offs*. Cambridge: Cambridge University Press.
- Kennedy, G., Berardo, A., Papavero, C., Horjus, P., Ballard, T., Dop, M. et al. (2012) Proxy measures of household food consumption for food security assessment and surveillance: comparison of the household dietary diversity and food consumption scores. *Public Health Nutrition*, 13, 2010–2018.
- Kirkpatrick, S. & Tarasuk, V. (2009) Food insecurity and participation in community food programs among low-income Toronto families. *Canadian Journal of Public Health*, Mar-Apr, 35–39.
- Kjaerulff, U.B. & Madsen, A.L. (2013) *Bayesian networks and influence diagrams: a guide to construction and analysis*. Berlin: Springer.
- Lagi, M., Bar-Yam, Y., Bertrand, K. & Bar-Yam, Y. (2012) Update February 2012. The food crises: a quantitative model of food prices including speculators and ethanol conversion. *arXiv:1203.1313*.
- Lee, A.M., Scharf, R.J., Filipp, S.L., Gurka, M.J. & DeBoer, M.D. (2019) Food insecurity is associated with prediabetes risk among U.S. adolescents, NHANES 2003–2014. *Metabolic Syndrome and Related Disorders*, 17, 347–354.
- Leonelli, M. & Smith, J.Q. (2015) Bayesian decision support for complex systems with many distributed experts. *Annals of Operations Research*, 1–26. <https://doi.org/10.1007/s10479-015-1957-7>
- Loopstra, R. (2014) Household food insecurity in Canada: towards an understanding of effective interventions. Ph.D. thesis, University of Toronto.

- Loopstra, R., Reeves, A. & Stuckler, D. (2015a) Rising food insecurity in Europe. *The Lancet*, 385(9982), 2041.
- Loopstra, R., Reeves, A., Taylor-Robinson, D., Barr, B., McKee, M. & Stuckler, D. (2015b) Austerity, sanctions, and the rise of food banks in the UK. *BMJ*, 350, h1775. <https://doi.org/10.1136/bmj.h1775>
- Pan, L., Sherry, B., Njai, R. & Blanck, H.M. (2012) Food insecurity is associated with obesity among US adults in 12 states. *Journal of the Academy of Nutrition and Dietetics*, 112, 1403–1409.
- Phillips, L.D. (1984) A theory of requisite decision models. *Acta Psychologica*, 56, 29–48.
- Pilgrim, A., Barker, M., Jackson, A., Ntani, G., Crozier, S., Inskip, H. et al. (2012) Does living in a food insecure household impact on the diets and body composition of young children? Findings from the Southampton Women's Survey. *Journal of Epidemiology and Community Health*, 66, e6.
- Queen, C.M. & Smith, J.Q. (1993) Multiregression dynamic models. *Journal of the Royal Statistical Society: Series B (Methodological)*, 55, 849–870.
- Raftery, A.E., Newton, M.A., Satagopan, J.M. & Krivitsky, P.N. (2007) Estimating the integrated likelihood via posterior simulation using the harmonic mean identity. *Bayesian Statistics*, 8, 371–416.
- Robert, C. & Casella, G. (2004) *Monte Carlo statistical methods*. New York: Springer.
- Seligman, H.K., Lاراia, B.A. & Kushel, M.B. (2010) Food insecurity is associated with chronic disease among low-income NHANES participants. *The Journal of Nutrition*, 140, 304–310.
- Smith, J.Q. (2010) *Bayesian decision analysis: principles and practice*. Cambridge: Cambridge University Press.
- Smith, J.Q., Barons, M.J. & Leonelli, M. (2015) Decision focused inference on networked probabilistic systems: with applications to food security. In: *JSM Proceedings 2015, Seattle, Washington: American Statistical Association*, pp. 3220–3233.
- Smith, J., Barons, M. & Leonelli, M. (2017) Coherent inference for integrating decision support systems. *arXiv preprint arXiv:1507.07394*.
- Tarasuk, V., Fitzpatrick, S. & Ward, H. (2010) Nutrition inequities in Canada. *Applied Physiology, Nutrition, and Metabolism*, 35, 172–179.
- Tarasuk, V., Mitchell, A. & Dachner, N. (2016) Household food insecurity in Canada, 2014. *Tech. rep.*, Toronto: Research to identify policy options to reduce food insecurity (PROOF). Available from: <http://proof.utoronto.ca/wp-content/uploads/2016/04/Household-Food-Insecurity-in-Canada-2014.pdf>
- Taylor-Robinson, D., Rougeaux, E., Harrison, D., Whitehead, M., Barr, B. & Pearce, A. (2013) The rise of food poverty in the UK. *BMJ*, 347, f7157. <https://doi.org/10.1136/bmj.f7157>
- The Economist Intelligence Unit. (2019) Global food security index. Web. Available from: <http://foodsecurityindex.eiu.com/>
- Tingay, R.S., Tan, C.J., Tan, N.C., Tang, S., Teoh, P.F., Wong, R. et al. (2003) Food insecurity and low income in an English inner city. *Journal of Public Health*, 25, 156–159. Available from: <http://jpubhealth.oxfordjournals.org/content/25/2/156.abstract>
- Trussell Trust. (2021) End of year stats April 2020–March 2021. *Technical report*, Trussell Trust.
- United Nations Office of the High Commissioner. (1966) International covenant on economic, social and cultural rights. <https://www.ohchr.org/en/professionalinterest/pages/cescr.aspx>, Last accessed on 11/02/2020.
- USDA. (2012) U.S. adult food security survey module. *Technical report*, USDA.
- West, M. & Harrison, J. (1997) *Bayesian forecasting and dynamic models*. Berlin: Springer.
- Wilkerson, R.L. & Smith, J.Q. (2021) Customised structural elicitation. *Expert Judgement in Risk and Decision Analysis*, 83–113. <https://doi.org/10.1080/01605682.2021.1907244>

How to cite this article: Barons, M.J., Fonseca, T.C.O., Davis, A. & Smith, J.Q. (2021) A decision support system for addressing food security in the United Kingdom. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 00, 1–24. <https://doi.org/10.1111/rssa.12771>

APPENDIX A. MEASURING THE ATTRIBUTES IN THE UTILITY FUNCTION

The utility function depends on the variables health, education, social unrest and costs which are defined as follows.

Health: Suppose the expert panellists define a proxy as a function of number of admission to hospital with diagnosis of malnutrition (primary or secondary) and number of deaths with malnutrition listed on the death certificate either as primary or secondary cause. Admissions data are available in the Hospital Episode Statistics (HES) from the UK government's Health and a Social Care Information Service which routinely links UK Office for National Statistics (ONS) mortality data to HES data. In the United Kingdom, the number of deaths caused primarily by malnutrition are very low and rates are not significantly different over time. Besides, malnutrition is usually accompanied by other diagnoses such as diseases of digestive system, cancers, dementia and Alzheimer's disease. Thus, the increase of deaths with malnutrition as a contributory factor might be due to ageing of the population and not due to food insecurity. Regarding admissions with malnutrition even the primary diagnosis numbers have increased over time with 391 in 2007–2008 and 780 in 2017–2018. Thus, in this work we considered the primary and secondary admission cases as a proxy for the health variable. Thus, the variable Health is defined as the count of finished admission episodes with a primary or secondary diagnosis of malnutrition coded ICD-10. An ICD-10 code of malnutrition on the episode indicates that the patient was diagnosed with, and would therefore being treated for malnutrition during the episode of care.

Education: The proxy for education could be defined as a function of educational attainment such as the proportion of pupils achieving expected grades in key stages 1, 2 and 4. Even though educational attainment is published annually at local and national levels by the UK government's Department for Education, the score system has changed in previous years and temporal comparisons are not adequate (Hill, 2014). Thus, as a proxy for education and its relation to food security we considered the proportion of pupils at the end of key stage 4 who were classified as disadvantaged. Thus, the variable Education is measured as the percentage of pupils at key stage 4 who were classified by the Department for Education as disadvantaged including pupils known to be eligible for free school meals (FSM) in any spring, autumn, summer, alternative provision or pupil referral unit census from year 6 to year 11 or are looked after children for at least one day or are adopted from care. Before 2015, this classification considered those who have been eligible for Free School Meals at any point in the last 6 years and Children who are 'Looked After'. In 2015, this definition was widened to also include those children who have been 'Adopted From Care'. Pupils classified as disadvantaged have a lower average educational attainment record than other pupils and there is a direct correlation between level of qualification and unemployment in later life; Poor educational attainment is strongly correlated with teenage pregnancy, offending behaviour, and alcohol and drug misuse. Comparisons between educational attainment for disadvantage and other pupils indicate a difference of 4.07 (2010/2011) and 3.66 (2016/2017) in the attainment gap index for Key stage 4 for state funded schools in England. The gap index is a score measuring the differences between the disadvantaged and non-disadvantaged groups in key level 2 and 4 (Hill, 2014). The index is the mean rank for all the disadvantaged and non-disadvantaged pupils divided by the number of pupils in each cohort. This decimal mean rank difference is scaled to 10 and ranges from 0 to 10, where a higher value means a higher attainment of non-disadvantaged compared to disadvantaged pupils. The index aims to be resilient to changes in the grading systems and in the assessments and curricula, and may be used for temporal comparisons.

Social Unrest: Inadequate food security can cause food riots (Lagi et al., 2012). In the United Kingdom, a riot is defined by Section 1(1) of the Public Order Act 1986 as where 12 or more persons who are present together use or threaten unlawful violence for a common purpose and the conduct of them (taken together) is such as would cause a person of reasonable firmness present at the scene to fear for his personal safety, each of the persons using unlawful violence for the common purpose is guilty of riot. Riot data are collected by the police. While the likelihood of a food riot is small in the United Kingdom currently, post-riot repairs both to physical environment and community relations can be considerable.

Costs: Costs of candidate intervention policies are routinely calculated and form part of the decision-making process. Indeed, as a response to falling budgets, DMs might revise the criteria for assistance of various kinds, for instance by making the eligible cohort smaller. Interventions which are effective but budget-neutral or cost-saving are obviously preferred, however, when the benefit of intervention may not be seen within the same financial year, this would form part of the decision-makers' discussion after the policies had been scored. This is the approach we take here, by scoring the policies and leaving the costs for final discussions of DMs.

APPENDIX B. DESCRIPTION AND DATA SOURCE FOR THE VARIABLES IN THE NETWORK

The variables Education and Health were described in Appendix A. The source for the Education data was <https://www.gov.uk/government/statistics> and for the health data was <https://digital.nhs.uk/data-and-information>.

Panel G1 (household income) is represented by the variable HIncome. This variable depends on the household income after expenses and is defined as follows:

- HIncome: Real net households adjusted disposable income per capita less the final consumption expenditure per head. Data source: <http://www.ons.gov.uk>

Panel G2 (food costs) is represented by the variable CFood and is defined as follows:

- CFood: CPI index of nine food groups, 2015 = 100. Food costs was measured by a combination of CPI indices of items representing household dietary diversity (Kennedy et al., 2012). The score is formed by nine food groups: cereals, meat, fish, eggs, milk, oils and fat, fruits, vegetables and beverages. Data source: <http://www.ons.gov.uk>

Panel G3 (income) accounts for access to credit (Lending), tax on the income (Tax), unemployment rate and social benefits and is defined as follows:

- Lending: Net lending (+)/net borrowing (-) by sector as a percentage of GDP - Household and non-profit institution serving households. Data source: <http://www.ons.gov.uk>
- Tax: Original household income minus post-tax income (deflated to 2018 index). Income has been equivalised using the modified-OECD scale. Data source: <http://www.ons.gov.uk>
- Unemployment: Male unemployment rate, aged 16 and over, seasonally adjusted. Data source: <http://www.ons.gov.uk>
- Benefits: Social assistance benefits in cash as a percentage of GDP. Data source: <http://www.ons.gov.uk>

Panel G4 (costs of living) accounts for expenditure per head (Living) and housing costs (Chousing) and is defined as follows:

- CLiving: Consumer price indices of the main variables composing the expenditures of a household: housing, including energy (CHousing), food (CFood), recreation (CRecreation), and transport (CTransport). Data source: <http://www.ons.gov.uk>
- CHousing: CPI of housing, water and fuels. Data source: <http://www.ons.gov.uk>

Panel G5 (food supply) accounts for output of food production (FProduction) and imports from European Union and other countries and is defined as follows:

- FProduction: Producer price inflation (Output of food products). Data source: <http://www.ons.gov.uk>
- FImports: Food imports from European Union countries plus imports from other countries. Data source: <http://www.ons.gov.uk>

Panel G6 (Oil costs) is represented by CPI of fuels and energy (COil and CEnergy) and is defined as follows:

- COil: Liquid fuels, vehicle fuels and lubricants (G) 2015 = 100. Data source: <http://www.ons.gov.uk>
- CEnergy: CPI of energy, 2015 = 100. Data source: <http://www.ons.gov.uk>

Panel G7 (Demography) is represented by part-time work rates (PartTime) and is defined as follows:

- PartTime: Part-time workers (Ill or disabled). Data source: <http://www.ons.gov.uk>

Panel G8 (Weather) is represented by number of days in which the air temperature falls below 0 degrees Celsius. In these cases, sensitive crops can be injured, with significant effects on production and is defined as follows:

- Frost: Number of days of air frost. Data source: <http://www.metoffice.gov.uk>

Panel G9 (Economy) accounts for economic context represented by Gross D domestic Product (GDP):

- GDP: Gross Domestic Product at market prices, seasonally adjusted. Data source: <http://www.ons.gov.uk>

APPENDIX C. R CODE AND DATA USED FOR ANALYSIS

The R codes and data used in the analysis of UK food security are freely available for download at the link: <https://github.com/thaiscofonseca/foodnetwork.git>